DELL DX Object Storage
File Gateway Deployment Guide

A Dell Technical White Paper

Dell | Storage

Storage Engineering

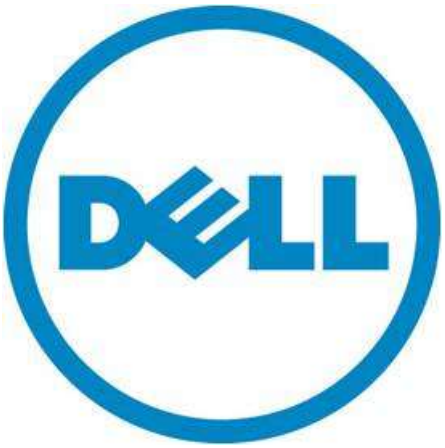

THIS DOCUMENT IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2011 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the *DELL* logo, and the *DELL* badge, and *PowerVault* are trademarks of Dell Inc. *Microsoft*, *Windows*, and *Active Directory* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. *Red Hat* and *Red Hat Enterprise Linux* are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

October 2011

Contents

Section I – Overview and Architecture 5
Scope of this Document and Requirements 5
Conventions Used in This Document 5
Working with Files and Commands 5
Deployment Checklist 6
Single-Server, Standalone Authentication, Local Spooler 6
Single-Server, Active Directory Member, Local Spooler 6
Failover, Standalone Authentication, External Spooler 6
Failover, Active Directory Member, External Spooler 7
Site Survey 7
Architecture 7
Standard hardware 7
Specific CFS Configurations 8
Single-server Configuration 8
Failover Configuration 9
Gateway Protocols 10
Other hardware 10
Software 11
How it all works together 12
DX Object Storage File Gateway (single-server) 12
DX Object Storage File Gateway (Failover) 13
See the *DX6000 User's guide* for Information on how to setup the DX6000 Cluster. 13
See the *DX Optimizer Node Deployment guide* for information on how to setup the DX6000G Optimizer Node. 13
Set up the DX6000G Cluster File Server 14
Assumptions and Requirements 14
Verify BIOS settings 14
Install and Configure Red Hat Enterprise Linux 15
Run Deployment Scripts 20
Set up Repositories and Packages 20
Configure Internal Spool/Cache (Single-Server Solution) 21
Configure External Spool/Cache (Failover Solution) 22
Configure Compression on a CFS Mount (Optional) 25
Configure Gateway Protocols 26

SMB/CIFS Gateway Service 26
Stand-alone Server (Workgroup Authentication) 27
Active Directory Domain Member Server 31
NFS Gateway Service 37
Before Configuring NFS 37
Configure NFS to share the CFS mount 38
Configure Share Resources 38
SMB/CIFS Shared Resource Configuration 38
Add a Share (Standalone Server) 38
Add a share (Active Directory Domain Member) 40
Remove a Share 41
NFS Shared Resource Configuration 42
Add an NFSv3 Share Point 42
Remove an NFS Share Point 42
Upgrading CFS 42
Administrative Maintenance Procedures 43
Starting CFS and CNS 43
Shut Down CFS and CNS 44
Special Considerations for MD3200i Spooler 44
Additional Information 45
Appendix A. Gateway Protocol Support 45
Protocol Gateway Limitations 45
Supported Protocols 45
Access Control Lists 45
SMB/CIFS Protocol Support 46
Appendix B. NFS Client Guidelines 46
Appendix C Manual Configuration Procedures 47
Create the Master Boot Record (MBR) on the Second Drive 47
Disable SELinux 48
Create the YUM Repository and Install Packages 49
Stop and Disable Services 51
Set up the NTP Server 51
Configure the Network Interfaces for Bonding 52
Configure the Domain Name Service Resolver 54
Install the CFS Software 55

Spooler and Cache File Systems 56
Create disk partitions and spooler file systems (Single-server solution) 56
External Spooler File System (Failover solution) 58
Configure the Cluster Name Space (CNS) 64
Configure the CFS and its DX Object Storage Mount Points 64

Section I – Overview and Architecture

Scope of this Document and Requirements

This document provides instruction for deploying either a CIFS or NFS gateway solution on the Dell DX Object Storage platform. Successful deployment enables customers to use a common file system with which they are comfortable.

The following is required before beginning the deployment:

- All DX Object Storage hardware has been racked and cabled.
- The Caringo CAStor software has been installed on the Cluster Services Node of the DX Object Storage system. (The CSN has the software factory installed by default.)
- The components of the solution are properly networked.
- You have all required network information (see Site Survey included in this document).

Conventions Used in This Document

The following fonts and conventions are used in this document to identify required actions and example text as it appears in the command line interface (CLI).

```
CLI text that you enter as part of a command or editing a file. (input)
Variable in code
Resulting text as it appears in a command-line interface. (output)
Emphasis in resulting text after a command is entered.
File name, directory name, or variable name
```

Working with Files and Commands

This document assumes you have a working knowledge of Linux. You can run the commands or navigate the file system either through a command line interface (CLI) or through the X-Windows GUI. In instances where CLI is the only option, the document will provide specific instruction on what you need to type.

Deployment Checklist

Complete the steps below for the type of CFS solution you are deploying.

Single-Server, Standalone Authentication, Local Spooler

(DX Cluster + 1 CFS, not a member of Active Directory)

- ☐ Complete the Site Survey
- ☐ Technical review.
- ☐ Order placement and delivery.
- ☐ Rack and cable (power and network) the DX Object Cluster and the CFS Server(s).
- ☐ Set up the DX Object Cluster.
- ☐ Set up the CFS Server.
- ☐ Configure SMB/CIFS gateway protocols for standalone server, or configure NFS gateway service.
- ☐ Add a SMB/CIFS or NFS share for standalone server.

Single-Server, Active Directory Member, Local Spooler

(DX Cluster + 1 CFS, member of Active Directory)

- ☐ Complete the Site Survey
- ☐ Technical review.
- ☐ Order placement and delivery.
- ☐ Rack and cable (power and network) the DX Object Cluster and the CFS Server(s).
- ☐ Set up the DX Object Cluster.
- ☐ Set up the CFS Server.
- ☐ Configure SMB/CIFS gateway protocols for Active Directory Member Server.
- ☐ Add a share (Active Directory Domain Member).

Failover, Standalone Authentication, External Spooler

(DX Cluster + 2 CFS Systems + MD3200i, not a member of Active Directory)

- ☐ Complete the Site Survey
- ☐ Technical review.
- ☐ Order placement and delivery.
- ☐ Rack and cable (power and network) the DX Object Cluster and the CFS Server(s).
- ☐ Set up the DX Object Cluster.
- ☐ Set up the CFS Server.
- ☐ Configure SMB/CIFS gateway protocols for standalone server, or configure NFS gateway service.
- ☐ Add a SMB/CIFS or NFS share for standalone server.
- ☐ Configure the Failover CFS Server.

Failover, Active Directory Member, External Spooler
(DX Cluster + 2 CFS Systems + MD3200i, member of Active Directory)

- ☐ Complete the Site Survey
- ☐ Technical review.
- ☐ Order placement and delivery.
- ☐ Rack and cable (power and network) the DX Object Cluster and the CFS Server(s).
- ☐ Set up the DX Object Cluster.
- ☐ Set up the CFS Server.
- ☐ Configure SMB/CIFS gateway protocols for Active Directory Member Server.
- ☐ Add a share (Active Directory Domain Member).
- ☐ Configure the Failover CFS Server.

Site Survey

Use the site survey to gather information about the customer site. You will use some of this information to determine which CFS solution is best for your customer, and much of the information is vital to a successful deployment.

DX CFS Site Survey.docx

Architecture

The DX Object Storage Gateway Solution is offered in two different types of configurations — single-server and failover.

Standard hardware

The following briefly describes the standard hardware components of DX Object Storage Gateway solutions. In most cases, none of these components are yet installed at the customer site.

DX6000G File Gateway CFS Server – This server hosts the content file server software that presents the Dell DX Object Storage system as a standard file system. In single-server solutions, this server also provides a spool cache on the local disk.

DX6000G Optimizer Node – This server hosts the Dell compression and decompression engines that are used to compress objects out-of-band and decompress objects in-line.

Cluster Services Node (CSN) – This node is an integrated services node that centralizes installation and configuration of both the network services required to run a DX Storage cluster and the software used to interface with it.

Storage Nodes – These nodes host data that is indirectly written to them by application servers (via the spool cache). A minimum of two nodes is required, and each data object has at least one replica.

Shared Storage System – Shared storage is a cache where files are stored before they get written to the DX Object Storage and also where they can be accessed on subsequent reads, if locally available. A separate shared storage system, such as the MD3200i is required for failover solutions. In these types of solutions, if the gateway server fails over to another gateway server, the common spool cache is still available to either write its data to the storage nodes, or serve data to clients.

Specific CFS Configurations

Two types of CFS configurations are offered – single server (no failover) and high availability (failover). A failover solution is recommended in environments that require high performance and/or have large file system configurations.

Single-server Configuration

The minimal CFS configuration is a single CFS server the uses internal storage for storing the spooled data and namespace. This configuration is recommended when CFS failover is not required and when the spooled data and namespace capacity requirements are less than 1TB.

Hardware

- (1) DX6000G server
  - Chassis: Four Hot-Plug Hard Drives, LCD diagnostics
  - Processor(s): 2 x E5620, 2.4GHz 12M Cache, Turbo, HT, 1066MHz Max Mem
  - Memory Configuration: 24GB Memory (6x4GB), 1066MHz Dual Ranked RDIMMs for Processors, Optimized
  - HDD Configuration: RAID1/RAID1 hard drive configuration
  - Hard Drives:
    - (2) 500GB 7.2K RPM Near-Line SAS 3.5in HDD for OS and application
    - (2) 1TB 7.2K RPM Near-Line SAS 3.5in HDD for spooled data and namespace
  - Primary Controller: SAS 6/iR SAS internal RAID adapter for Hot Plug Configuration, PCI-Express
  - PSU: 500Watt Redundant power supplies
  - Embedded Management: iDRAC6 Enterprise
  - Network Adapter: Broadcom NetXtreme II 5709 Gigabit NIC w/TOE & iSOE, Quad Port, Copper, PCIe-4
  - DVD-ROM Drive
  - Operating System: RHEL 6 X64 Enterprise with 3 Year Subscription

Network

The DX6000G CFS server is configured with one Broadcom dual port BCM 5716 controller and one Broadcom quad port NetExtreme II 5709 GbE controller for a total of six GbE ports. The recommended network port allocation for a single-server configuration is:

Connect NIC ports 0 and 1 to the CIFS/NFS network.

Connect NIC ports 2 thru 5 to the DX Storage Cluster network.

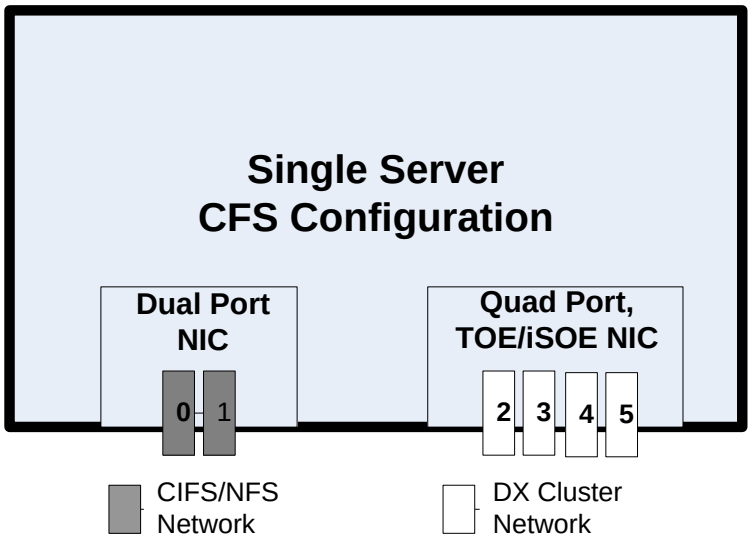

Failover Configuration

The failover configuration allows DX6000G CFS servers to be deployed in pairs so that if one server fails, the share can be recovered by the other CFS to minimize recovery time. In this configuration, the CFS servers utilize iSCSI storage to store spooled data.

Hardware

This configuration consists of the following hardware:

- (2) DX6000G Cluster File Servers
  - Chassis: Up to four Hot-Plug Hard Drives, LCD diagnostics
  - Processor(s): 2 x E5620, 2.4GHz 12M Cache, Turbo, HT, 1066MHz Max Mem
  - Memory Configuration: 24GB Memory (6x4GB), 1066MHz Dual Ranked RDIMMs for Processors, Optimized
  - HDD Configuration: RAID 1 configuration; internal drives in RAID 1 for OS and application; spooled data and namespace stored on iSCSI
  - Primary Controller: SAS 6.iR SAS internal RAID adapter for Hot Plug Configuration, PCI-Express
  - Hard Drives: (2) 500GB 7.2K RPM Near-Line SAS 3.5in HDD
  - PSU: 500Watt Redundant power supplies
  - Embedded Management: iDRAC6 Enterprise
  - Network Adapter: Broadcom NetXtreme II 5709 Gigabit NIC w/TOE & iSOE, Quad Port, Copper, PCIe-4
  - DVD-ROM Drive
  - Operating System: RHEL 6 X64 Basic with 3 Year Subscription

- (1) iSCSI Storage: MD3200i
  - Controllers: Dual controller option
  - Hard Drives: Up to 12 3.5-inch HDD, 500GB near-line SAS 6GB, 7.2K, 3.5in HDD
  - HDD Configuration: RAID5

Network

For a failover configuration, the iSCSI storage can be on its own switch isolated from the DX Storage Cluster or it can be attached to the DX Cluster network. In cases where iSCSI storage is attached to the DX Cluster network, ports 2 thru 5 can be shared between the iSCSI storage and DX Storage Cluster. In this configuration, Dell recommends configuring network ports for adaptive load balancing, but sites

that have suitable competence may prefer configuration and management of link aggregation control protocol in place of adaptive load balancing.

The preferred failover network configuration is identical to a single-server CFS configuration (above), except that iSCSI traffic shares the four NIC ports on the private DX Storage network. If a site wishes to use pre-existing iSCSI storage available on its own VLAN, the recommended network port allocation is:

- Connect NIC ports 0 and 1 to the CIFS/NFS network.
- Connect NIC ports 2 and 3 to the DX Storage Cluster network.
- Connect NIC ports 4 and 5 to the iSCSI Storage network.

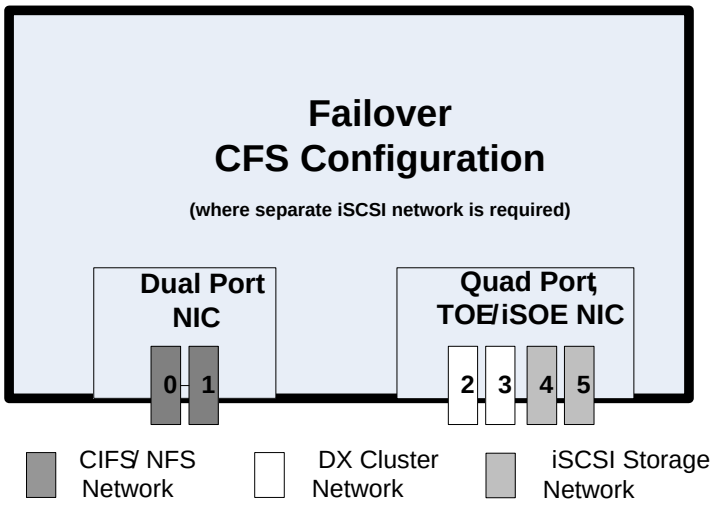

Gateway Protocols
The DX6000G CFS can be configured for NFS or SMB/CIFS gateways. For SMB/CIFS, users and groups can access the system through standalone authentication, or if Active Directory Services exists, users and groups can be authenticated through the existing ADS structure. For NFS, Dell supports version 3 only; it does not support version 4.

Other hardware

The following briefly describes the hardware components required to work as part of DX Object Storage Gateway solutions. They may already be present in the customer environment, or ordered as part of the solution.

Network Switch(es) – Configuring the solution requires extensive knowledge of how the customer uses switches to segment VLANs. Some customers may run each VLAN through a separate switch, and some may segment a switch for multiple VLANs.

Application Server(s) – These servers write data to the spool cache storage in the Gateway solution. In a non-gateway DX Object Storage solution, they write directly to the storage nodes.

Domain Controller – This server manages logins, authentication, groups and permissions. Domain controller information is an essential part of setting up a gateway solution.

Software

The following briefly describes the standard software components for DX Object Storage Gateway solutions.

Red Hat Enterprise Linux – This is the operating system residing on the Cluster Services Node and the CFS Server. Different versions may run on the DX6000G CFS and the CSN. See the interoperability matrix for information about versions.

CFS – This is the software on the DX6000G CFS that presents Dell DX Object Storage to clients as a common file system.

Cluster Services – These services reside on the CSN and enable you to configure the CSN, define user access, set parameters for backup and restore, and define SCSP Proxy settings.

Content Storage – These services reside on the CSN and enable you to define properties, licensing, and user access for the whole cluster.

Content Router – These services provide replication of content between remote clusters and enumeration of DX Object Storage content for other purposes like search indexing or virus scanning. Content Router is not required if the cluster is not replicating to a remote cluster.

How it all works together

The data flow in a DX Object Storage File Gateway depends on the type of configuration. The following examples showa DX Object Storage cluster with a single-server gateway, and a DX Object Storage cluster with a failover gateway.

DX Object Storage File Gateway (single-server)

In a DX Object Storage File Gateway, data objects are written to a DX6000G CFS server before being written to a Storage Node. In this configuration, the application server and clients are actually viewing objects as they reside in the spool cache of the CFS. These objects are presented to the end user as part of a file system.

Figure 1. DX Object Storage File Gateway (single-server)

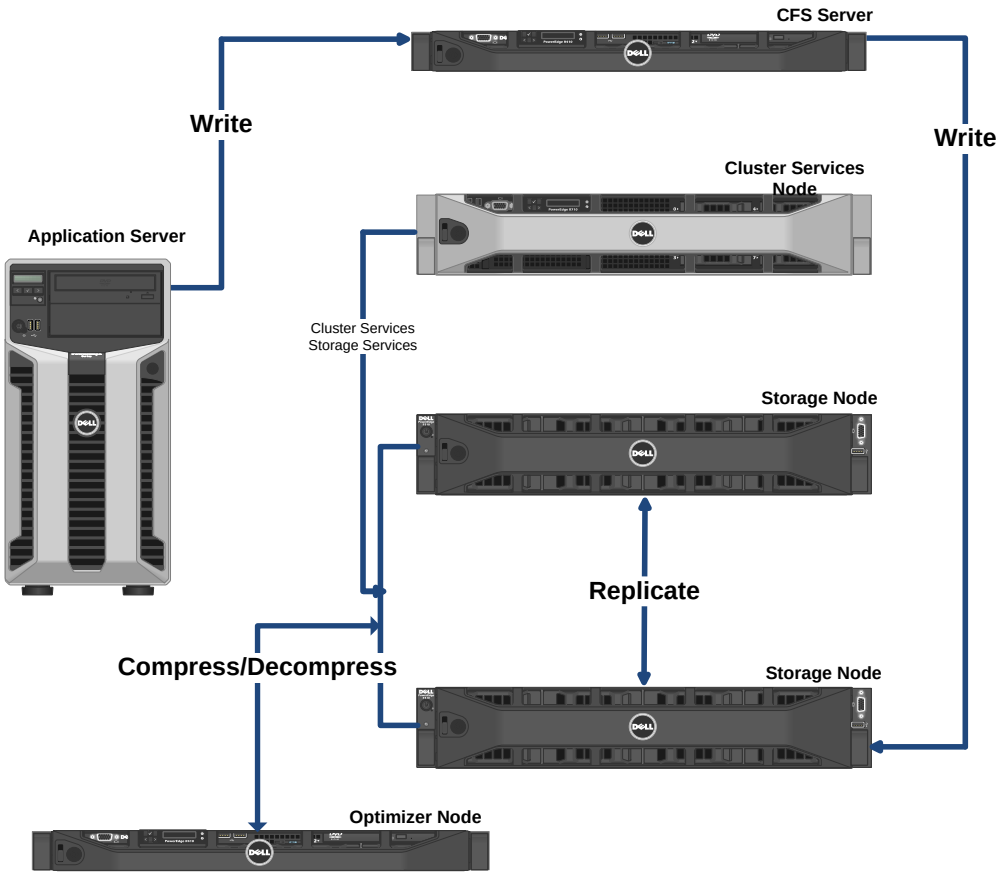

DX Object Storage File Gateway (Failover)

A failover configuration (see Figure 2) provides two DX6000G CFS servers and a separate dedicated spool/cache. This configuration provides continuous service on the gateway, as long as the cluster and the shared storage are running.

Figure 2. DX Object Storage File Gateway (Failover)

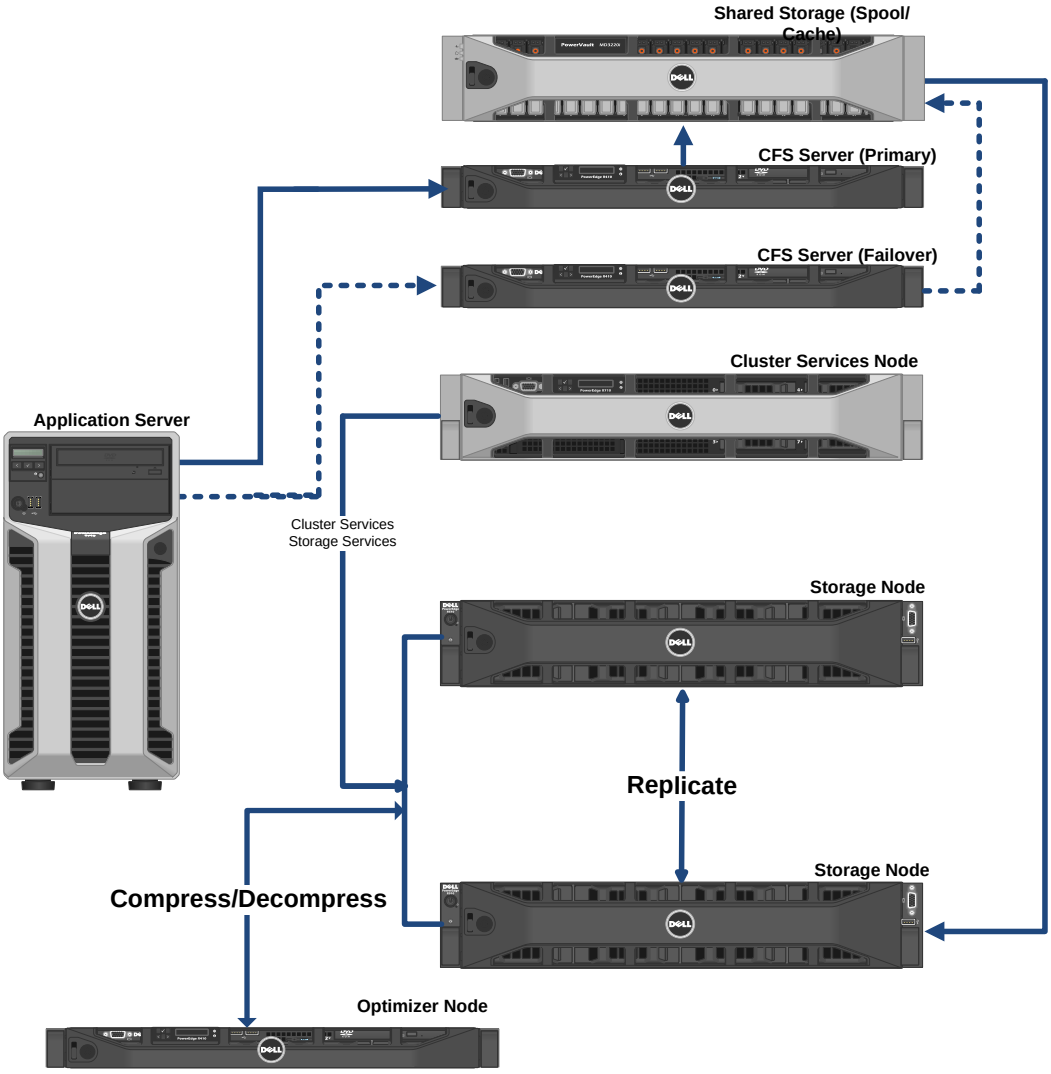

See the *DX6000 User's guide* for Information on how to set up a DX Object Storage cluster.

See the *DX Optimizer Node Deployment guide* for information on how to setup the DX6000G Optimizer Node.

Set up the DX6000G Cluster File Server

BEFORE YOU BEGIN:

- Did you complete the Site Survey?
- Did you Set up the DX Object Storage Cluster? See the *DX Object Storage Platform User's Guide* and the *DX Object Storage Cluster Services Node Installation and Configuration Guide*.

***NOTE**: Make sure there is a DNS entry for the CIFS/NFS interface of the server. If the site does not have a DNS server, make sure that hostname is resolvable to the CIFS/NFS interface IP address from the /etc/hosts file.* This section includes the steps for setting up and activating the CFS. The procedures should be performed in the following order:

- Verify the necessary BIOS settings.
- Configure the operating system.
- Run the installation scripts.

Assumptions and Requirements

- ☐ Red Hat Enterprise Linux version 6 or later is factory-installed on the CFS.
- ☐ You have configured the CSN, and the cluster is up and running.

***IMPORTANT**: Before you begin the installation, ensure that you have the following information, which will be required during the installation of Red Hat and the CFS software:*

- ☐ Public network address
- ☐ Public network subnet
- ☐ Private network address
- ☐ Private network subnet
- ☐ Public network DNS
- ☐ Public network gateway
- ☐ Storage Node IP addresses
- ☐ iSCSI storage device IP address (for failover solutions)

Verify BIOS settings

For the CFS to function, several options must be set in the BIOS <F2>:

- ☐ All processors and cores enabled
- ☐ In the Integrated Devices category, Gb NICs enabled (not with PXE)

NEXT STEP: Install and Configure Red Hat Enterprise Linux

Install and Configure Red Hat Enterprise Linux

The DX6000G CFS has Red Hat Linux factory-installed. Follow the steps below to re-install the Red Hat Linux Operating System with the configuration parameters required for DX6000G CFS. Before powering on the system, ensure you have an external connection.

1. Power on the DX6000G system and insert the Red Hat Enterprise Linux x64 DVD.
2. When the first Red Hat screen appears, click **Next**.
3. Select the installation language and click **Next**.
4. Select the appropriate keyboard and click **Next**.
5. Select **Basic Storage Devices** as the type that will be part of your installation and click **Next**.
6. When a screen appears stating that a previous version has been detected, select **Fresh Installation**.
7. Enter a fully-qualified **Hostname** for the CFS that will identify it on the network (for example, server.domain.tld).

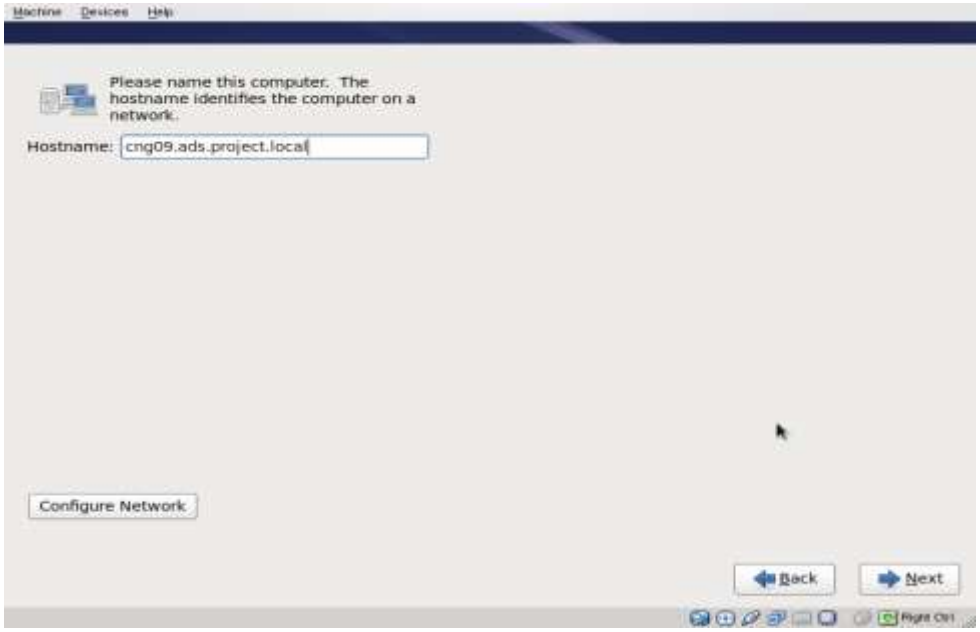

NOTE: *This hostname also needs to be added to the Domain Name Server itself.*

8. Click **Configure Network**.
9. On the **Wired** tab of the **Network Connections** screen, select **System eth0**, and click **Edit**.
10. Enter a **Connection name** and ensure that the **Connect automatically** box is checked and **MTU** is set to **automatic**.
11. Click the **IPv4 Settings** tab, and then click **Add** and select **Manual** as the **Method**.
12. Enter the **Address**, **Netmask**, and **Gateway** for the CFS (see Site Survey).

13. Provide information about the **DNS servers** and **Search domains** where the CFS server will reside.

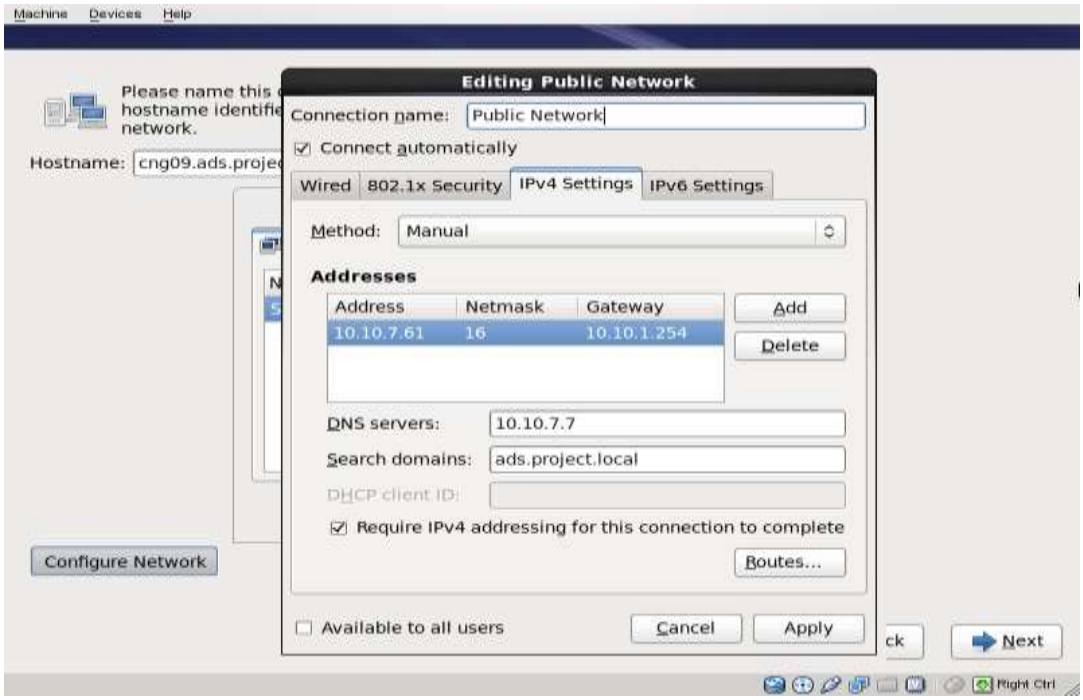

14. Click **Apply**, click **Close** to exit the **Network Connections** screen, and then click **Next**.

15. In the time zone screen, select the city in your time zone, check **System clock uses UTC** and click **Next**.

16. On the password screen, enter the **Root Password**, enter it again in **Confirm**, and then click **Next**.

17. On the **Installation Type** screen, select **Use All Space**, check **Review and modify partitioning** layout, and click **Next**.

Machine Devices Help
Which type of installation would you like?
Use All Space
Removes all partitions on the selected device(s). This includes partitions created by other operating systems.
Tip: This option will remove data from the selected device(s). Make sure you have backups.
Replace Existing Linux System(s)
Removes only Linux partitions (created from a previous Linux installation). This does not remove other partitions you may have on your storage device(s) (such as VFAT or FAT32).
Tip: This option will remove data from the selected device(s). Make sure you have backups.
Shrink Current System
Shrinks existing partitions to create free space for the default layout.
Use Free Space
Retains your current data and partitions and uses only the unpartitioned space on the selected device (s), assuming you have enough free space available.
Create Custom Layout
Manually create your own custom layout on the selected device(s) using our partitioning tool.
Encrypt system
Review and modify partitioning layout
Back
Next

18. In the storage devices screen, select *only the boot drive* (smallest drive) in the **Data Storage Devices** list, and click the arrow to move it to the **Install Target Devices** list and click **Next**.

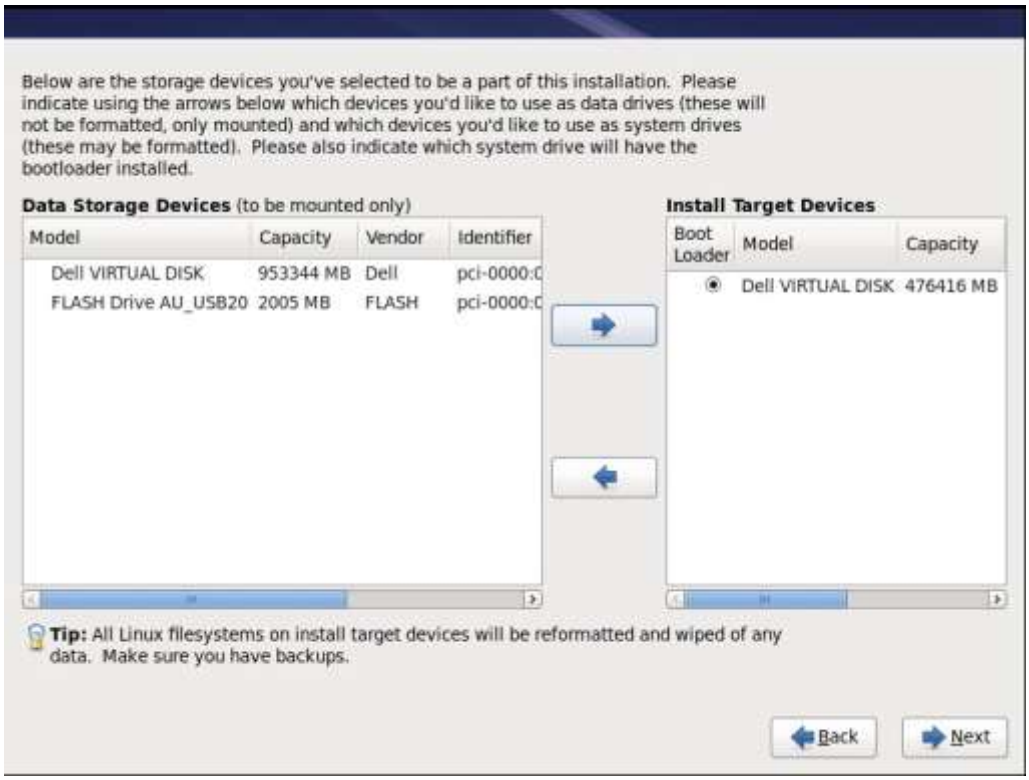
Below are the storage devices you've selected to be a part of this installation. Please indicate using the arrows below which devices you'd like to use as data drives (these will not be formatted, only mounted) and which devices you'd like to use as system drives (these may be formatted). Please also indicate which system drive will have the bootloader installed.
Data Storage Devices (to be mounted only)
Model Capacity Vendor Identifier
Dell VIRTUAL DISK 953344 MB Dell pci-0000:0
FLASH Drive AU_USB20 2005 MB FLASH pci-0000:0
Install Target Devices
Boot Loader Model Capacity
Dell VIRTUAL DISK 476416 MB
Tip: All Linux filesystems on install target devices will be reformatted and wiped of any data. Make sure you have backups.
Back
Next

19. In the **Please Select a Device** screen, delete all partition layouts and configurations.

20. Edit the partition layout as follows:

Note: In a single-server configuration, where there are two RAID 1 groups, the 1st (OS) RAID group may show up as sdb

a. Create the boot partition.
    i. Select the free space under sda.
    ii. Click **Create**.
    iii. Select **Standard Partition** and click **Create**.
    iv. Enter **/boot** for **Mount Point**.
    v. Ensure that **sda** is selected under allowable drives, and enter 1024 for **Size**.
    vi. Check **Force to be a primary partition** and click **OK**.
b. Create the logical volume for the root.
    i. Click **Free** under **sda1** and click **Create**.
    ii. Select **LVM Physical Volume** and click **Create**.
    iii. Ensure that **sda** is the only drive selected, and select **Fill to maximum allowable size**. And click OK.
    iv. Select the physical volume you just created and select **LVM Volume Group** and select **Create**.
    v. In the **Make LVM Volume Groups** screen, click **Add** to make a logical volume.
    vi. Enter **/** for the **Mount Point**, **ext4** for **File System Type**, and **LV_root** for **Logical Volume Name**.
    vii. Enter **51 GB** for **Size** and click **OK**.
c. Configure the remainder of disk (sda2) as an LVM physical disk (primary partition), using the following values:
    - Lv_var (**Volume Name**), 25G (**Size**) /var (**Mount Point**) ext4 (**File system type**)
    - Lv_swap (**Volume Name**), 24G (**Size**) swap (**File system type**)
    - Leave the remainder of LVM space unallocated and available for customers to create new volumes as needed or expand existing volumes.

21. Click **Next** to initiate formatting the hard drives, and click **Format** when asked to confirm.

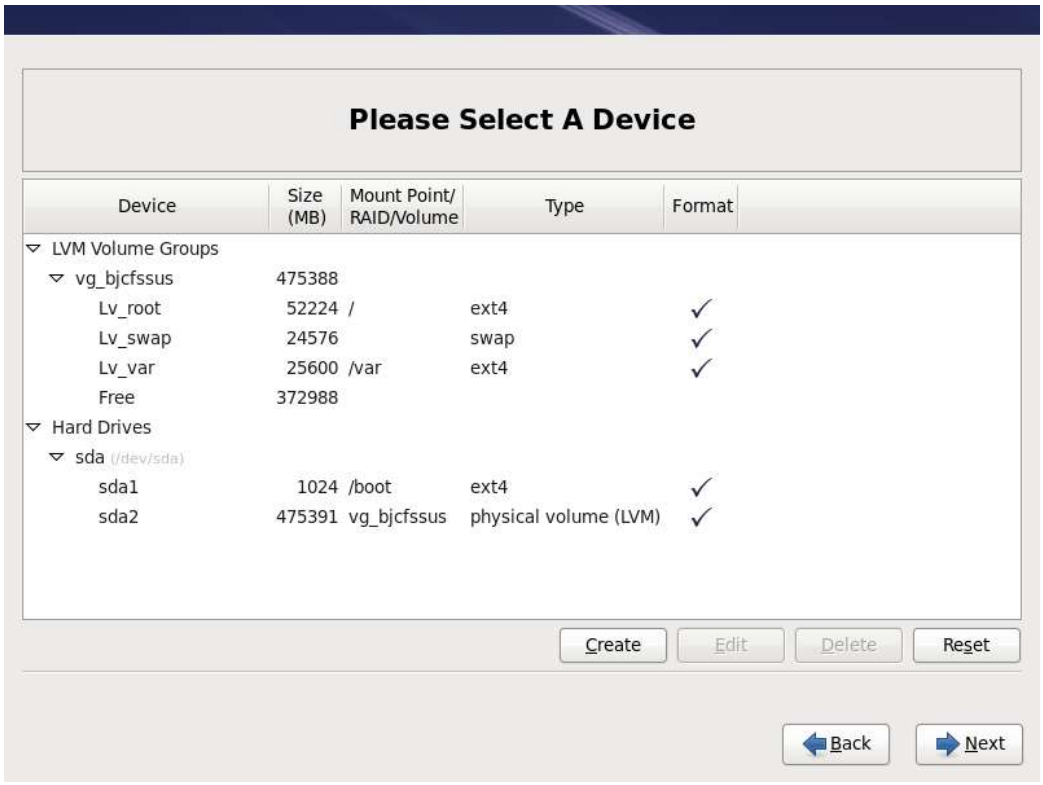

22. Click **Write changes to disk** when prompted.
23. When the boot loader operating system list appears, click **Next**.
24. Select **Desktop** as the installation type and click **Next**.

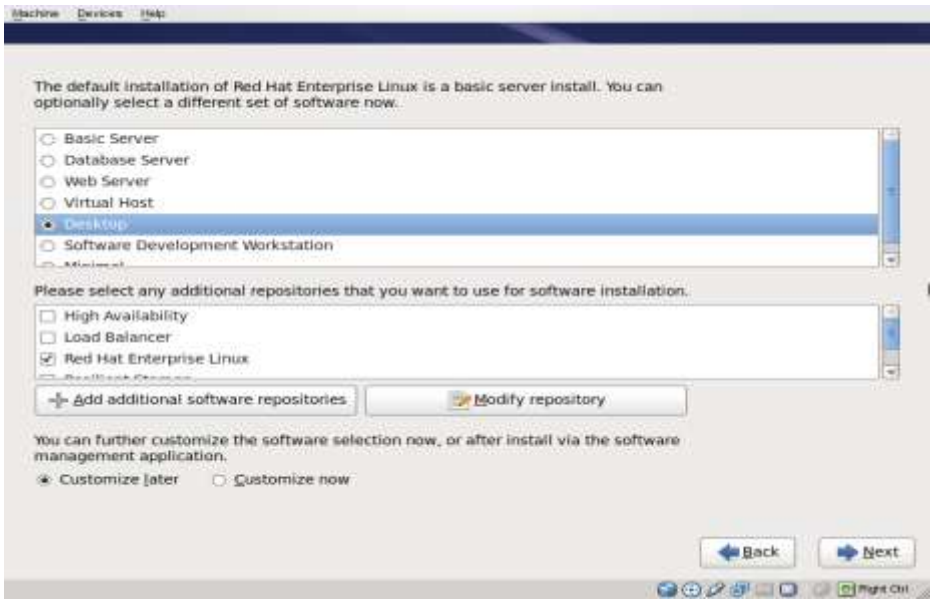
The default installation of Red Hat Enterprise Linux is a basic server install. You can optionally select a different set of software now.
Basic Server
Database Server
Web Server
Virtual Host
Desktop
Software Development Workstation
Please select any additional repositories that you want to use for software installation.
High Availability
Load Balancer
Red Hat Enterprise Linux
Add additional software repositories
Modify repository
You can further customize the software selection now, or after install via the software management application.
Customize later
Customize now
Back
Next

25. Click **Next** to start the installation.
26. After the installation completes, click **Reboot** when prompted.

27. When the **Welcome** screen appears, click **Forward** to advance through the **License** and **Software Update** screens.

28. In the **Create User** screen, enter a **Username**, **Full Name**, and **Password**, and click **Forward**.

Machine Devices Help
Welcome
License Information
Set Up Software Updates
Create User
Date and Time
Kdump
Create User
You must create a 'username' for regular (non-administrative) use of your system. To create a system 'username', please provide the information requested below.
Username: sysadmin
Full Name: System Admin
Password:
Confirm Password:
If you need to use network authentication, such as Kerberos or NIS, please click the Use Network Login button.
Use Network Login...
If you need more control when creating the user (specifying home directory, and/or UID), please click the Advanced button.
Advanced...
Back
Forward

29. On the **Date and Time** screen, select the **Synchronize data and time over the network** checkbox and click **Forward**.

30. On the **Kdump** screen uncheck the box.

Run Deployment Scripts

After the operating system has been installed, you will run two deployment scripts that complete the Dell DX6000G Cluster File Server (CFS) configuration.

Set up Repositories and Packages

1. Download the Dell DX CFS Software package from iDrive or support.dell.com.

2. Ensure that the RHEL DVD is in the drive.

3. Make sure you are logged in as the root user.

4. Copy **the Dell-DX-CFS-Software-2.6.4.zip** file to a directory e.g cp Dell-DX-CFS-Software-2.6.4.zip /home/administrator/Desktop

5. Extract the zip file and cd to the cfs_install folder.

6. In terminal mode, navigate to the folder where the scripts are extracted, type **`./phase1.sh`** and press <Enter>.

   This part of the script creates a master boot record on all scsi devices, which ensures that you can reboot the system if a drive fails.

7. When prompted, press <Enter>.

This copies the RHEL media to local drive to configure the YUM repository, installs local package dependencies, disables SE Linux and all the services, as necessary

8. When prompted, select whether internal or external storage will be used and press <Enter>.

NEXT STEPS: Configure Internal Spool/Cache (Single Server) **OR** Configure External Spool/Cache (Failover Solution)

Configure Internal Spool/Cache (Single-Server Solution)

1. When prompted for the bonding mode of the public network, select **balanced-Alb (Adaptive Load Balancing)** or **LACP**.
2. Enter the IP address for the public network (ports 0-1).
3. Enter the netmask for the public network.
4. Enter the gateway IP address.
5. Enter the primary DNS server address for the Public network, or leave empty if no DNS server is used.
6. Enter the domain name for the public network.
7. Select the bonding mode for the private network (ports 2-5). (Typically, this would be Balanced-Alb).
8. Enter the IP address for the private network (ports 2-5)
9. Enter the netmask for the private network.
10. Enter the gateway IP address for the private network or leave empty if no gateway is used.
11. Enter the DNS server for the private network, or leave empty if no DNS is used.
12. Enter a domain name for the private network, or leave empty if none is used.
13. Verify that network information is correct and press Enter.

    The bonding completes and CFS is installed. At the end of the script, a message states that Phase 1 is complete.

14. Eject the DVD and reboot.
15. Navigate to the folder where the scripts are extracted and type **`./phase2.sh`**.
16. Enter the name of a CFS spooler volume to be created and press <Enter>.
17. Create any more CFS spooler volumes, and then select No and press <Enter> once all volumes entered.

18. When prompted, select **Yes** to use ZeroConf to locate the DX Object Storage, or select **No** and enter the IP address of the primary DX Object Storage node.

Note: The CFS Gateway must be connected to the DX private network.

19. If you entered an IP address for the primary access node, enter an IP address for the secondary access node.
20. Enter the percentage of internal storage you want to remain unused. This is the percentage of storage that will not be part of any spoolers and used future expansion.
21. Enter the percentage of storage you want to use for the CNS cache
22. Enter the percentage of storage you want to use for each CFS spooler
23. Review the summary and if the percentages are OK, select **Yes**
24. When the root anchor UUID displays, document the UUID and store it in a safe place.

    This CNS root anchor UUID is required if you ever need to recover the CFS.

25. Start CNS when prompted.

    A message states the CNS configuration is complete.

26. When a message states the spool/cache directory does not exist, select **Yes** to create one.
27. When a message states the a mount does not exist, select **Yes** to create one.
28. When prompted, select **Yes** to mount the directory.
29. Repeat for other CFS volumes.

When a message in the script states **All Done**, the configuration is complete and you are ready to configure the gateway protocol.

Configure External Spool/Cache (Failover Solution)

1. Select whether iSCSI will be on its own network.

   If you select Yes, you will define the following for public network, private network, and iSCSI network. If you select No, you define the following only for public and private networks.

   The following questions (for the public network) are examples of the questions you will be asked about each network:

   - Select **balanced-Alb (Adaptive Load Balancing)** or **LACP (Link Aggregation Control Protocol)**.
   - Enter the IP address.
   - Enter the netmask.
   - Enter the gateway IP address.
   - Enter the primary DNS server address, or leave empty if no DNS server is used.
   - Enter the domain name.

2. Verify the network information is correct and press Enter.

   The bonding completes and CFS is installed. At the end of the script, a message states that Phase 1 is complete.

3. Eject the DVD and reboot.

4. Navigate to the folder where the scripts are extracted and type **`./phase2.sh`**.

5. Enter the name of a CFS volume to be created and press <Enter>.

6. Create any more CFS volumes, and then select No and press <Enter> once all volumes entered.

7. When prompted, select **Yes** to use ZeroConf to locate the DX Object Storage, or select **No** and enter the IP address of the primary DX Object Storage node.

8. If you entered an IP address for the primary access node, enter an IP address for the secondary access node.

9. Enter the IP address for the iSCSI device.

10. Enter an iSCSI Qualified Name (IQN) for the host.

11. When a message appears showing an IP address for the storage device, select **Yes** and press <Enter>.

12. Configure your storage device so volumes are available to the host. For details, refer to the *Dell™ PowerVault™ MD3200i Deployment Guide* available on support.dell.com. The following steps may be used as an example

Configure the MD3200i

The MD3200i spooler supports a minimum of six drives and a maximum of 20. The drives should be configured as RAID5 with one hot-spare. You will also need to create a number of LUNs (logical disks). The number of LUNs you should create is based on the number of CFS shares, maximum file size, and performance requirements. You will need one LUN for the CNS cache and one for each CFS file system.

1. Cable the MD3200i as follows:

   a. Connect the management ports on each controller to the public network.

   b. Connect the data ports (4 on each controller) to the storage network.

2. Install the PowerVault Modular Disk Storage software (also known as MDCU) on a Windows or Linux management station.

   IMPORTANT: *Do not install any Dell PowerVault Linux drivers.*

   NOTE: *The management station should be separate from the CFS or CSN servers.*

3. Run the management software and allow it to automatically discover storage devices.

   NOTE: *Autodiscovery assumes that the management station is on the same subnet of the public network. If the Autodiscovery does not begin automatically when launching the*

application, select Tools→Autodiscovery in the console.

4. Configure the disk group.

    a. Open the MD3200i management screen by double-clicking on storage array.

    b. Click the **Logical** tab, right-click a disk, and click **Create** to configure the drives on the MD3200i into a single RAID5 array, leaving one disk as hotspare.

    c. Click the **Logical** tab, right-click an array, and click **Create** to configure the LUNs.

    NOTE: *Use essentially the same logic as you did for sizing the storage nodes on a DX cluster when creating virtual disks in the shared spool. For example, if the customer has larger file sizes, you should create larger LUNs. If you have reason to believe that the customer may add additional CFS file systems in the future, you can create extra LUNs for them to use when necessary, or leave unconfigured space for expansion.*

5. Create a host group on the MD3200i for each failover node pair.
    a. Mappings menu → Define → Host Group
    b. Enter host group name and click OK.

6. Assign the LUNs you created to the host group with whatever LUN numbers are desired.

    a. Expand Undefined Mappings, right click on the LUN and then select Define Additional Mappings.

    b. Select Host group or host.

    c. Select LUN # and click Add.

7. Configure iSCSI on the MD3200i.

    a. Setup tab.

    b. Configure iSCSI Host Ports

    c. IP address and subnet mask

    d. Don't need gateway

    e. Select iSCSI host ports for other data ports and configure the same.

    f. If VLAN, click Advanced IPv4 Settings and enter VLAN information.

    g. Click OK.

    Check Manage iSCSI Settings and Target Authentication set to None.

8. Make the LUNs visible to the host:
    a. Mappings tab.
    b. Select the host group with the LUNs, and select **Define→Host**.
    c. Enter a **User label**.
    d. Select **Add by selecting a known unassociated host port identifier**.
    e. Select an entry from the **known unassociated host port identifier** drop-down list.
    NOTE: *Click the Refresh button if a host port does not initially appear in the list.*
    f. Enter a host name and Click **Add**, and click **Next**
    Select **Linux** and then **Finish**.

This allows the host to see the LUNs; prior to this step all it could see was an 'access volume', which is used for in-band management of the storage.

You should have at least as many volumes available as the number of CFS volumes you created, plus one.

13. Once the storage device has been configured and the storage volumes are created, select **Yes** and press <Enter>.
14. Select the volumes and press **OK**.

    Partitions are created on the iSCSI LUNs and formatted into an ext 4 file system. CNS admin then runs to configure Caringo Name Space.

15. When the root anchor UUID displays, document the UUID and store it in a safe place.

    NOTE: *This UUID is required if you ever need to recover the CFS. Ensure you save this UUID in a safe place.*

16. Start CNS when prompted.

    A message states that CNS configuration is complete.

17. When a message states that spool/cache directory does not exist, select **Yes** to create one.
18. When a message states that a mount does not exist, select **Yes** to create one.
19. When prompted, select **Yes** to mount the directory.
20. Repeat for other CFS volumes.

    When a message in the script states **All Done**, the configuration is complete and you are ready to configure the gateway protocol.

Configure Compression on a CFS Mount (Optional)

BEFORE YOU BEGIN: See the *DX Optimizer Node Deployment Guide* for setup information.

The Dell DX6000G Gateway v2.6.4 supports file compression by default. Objects are compressed by Dell DX Object Storage Compression Software that runs on a separate DX6000G Optimizer node. The DX6000G Gateway writes files as mutable objects to the DX Storage Cluster, and compression is enabled at the CFS mount level so that all subsequent new files and existing file revisions written to the mount are compressed as a background process by the DX6000G Optimizer Node. Compressed files that are read from the mount, are decompressed in-line by the DX6000G Optimizer node.

After completing phase 2 of the deployment scripts, set compression on the mount point by running a command, such as the following:

```
cfs-admin policy --add --reps=2 --del=yes --span=6m --compress=fast /mnt/MyCFSMount
```

This command creates a mount to MyCFSMount. The policy defines states that after 6 months, the object will be compressed with fast compression, and have 2 replicas that will be deletable. (The other compression option is *best*, which is more compression but takes longer.)

NOTE: Metadata and data updates to mutable objects as a result of compression are not registered by the Content Name Space (CNS). Therefore, a compressed file will be visible from the DX6000G Gateway as if it were the original uncompressed file.

NEXT STEP: Configure Gateway Protocols

Configure Gateway Protocols

In addition to being able to write to a locally mounted Linux file system, the CFS platform design makes it possible to layer network file services over the Dell DX Object Storage mounted file system using any software that makes basic operating system calls to access a file system.

The CFS implements a SMB/CIFS protocol gateway to the Dell DX Object Storage platform, and also an NFS protocol gateway. This section explains how these two gateway services can be configured in two stages:

1) Configuring the protocol gateway service

2) Adding CIFS/NFS shared storage resources.

SMB/CIFS Gateway Service

CFS SMB/CIFS protocol gateway services can be configured either manually or using the CFS-admin cifs-server utility. The CFS-admin tool can be used to configure the CFS server either as a stand-alone server that performs purely local authentication, or as a member of a Microsoft Active Directory security domain.

***NOTE**: A customer's gateway can be configured only as standalone (local authentication) server OR as an Active Directory Domain member server. It cannot be configured as both; it must be one or the other.*

Where configured as an Active Directory (AD) member, you can set file and directory access permissions using Microsoft Windows ACLs. This requires support for POSIX ACLs in the underlying file system.

The procedures outlined in this section cover only the configuration of the mode of service that the SMB/CIFS protocol gateway will provide. Configure shares so that Microsoft Windows workstations and servers can access Dell DX Object Storage resources.

Stand-alone Server (Workgroup Authentication)

***NOTE**: A customer's gateway can be configured only as standalone (local authentication) server OR as an Active Directory Domain member server (see* Active Directory Domain Member Server*). It cannot be configured as both; it must be one or the other.*

Microsoft Windows SMB/CIFS networking makes heavy use of name-to-IP address resolution methods. The older methods use NetBIOS (Network Basic Input/Output System) over TCP/IP technologies and depend either on UDP broadcasts-based name resolution processes, or use WINS (Windows Internetworking Name Service). Newer methods depend on DNS.

NOTE: *Where the CFS SMB/CIFS server is configured to operate as a standalone server (i.e.: makes use of local authentication)* ***it is highly recommended to use both WINS and DNS****, but at least one of these (WINS or DNS) must be correctly configured.*

1. Run the following command to save the original */etc/samba/smb.conf* file.

```
#mv /etc/samba/smb.conf /etc/samba/smb.conf.orig
```

2. Run the following command to open the */etc/samba/smb.conf* file.

```
#vi /etc/samba/smb.conf
```

3. Replace the *workgroup* name *MYGROUP* and the *netbios name* names (in upper case characters – each max 14 characters) that are appropriate for the site:

```
[global]
        workgroup = MYGROUP
        netbios name = CIFSFS
        server string = DX Storage
        log level = 1
        log file = /var/log/samba/log.%L.%m
        max log size = 0
        load printers = No
        disable spoolss = Yes
        os level = 0
        posix locking = No
```

NOTE: *If the site uses a WINS server, add the following to the above:*

```
        wins server = 123.45.67.89
```

(where 123.45.67.89 should be replaced with the IP address of the WINS server for the site)

3. Start the Samba daemons in preparation for the final CFS resource configuration.

a. From a *root* login shell, run these commands to set *smbd* and *nmbd* to start automatically at boot time:

```
# chkconfig smb on
# chkconfig nmb on
```

b. Start the server daemons by running the following commands:

```
# service nmb start
# service smb start
```

c. Verify that the daemons are running as shown here:

```
# ps ax | grep mbd
8099  ?        Ss     0:00 smbd -D
8113  ?        Ss     0:01 nmbd -D
8139  ?        S      0:00 smbd -D
...
```

4. Create an administrative account for the local SMB/CIFS server, using either the root account (easiest) or a normal user account.

   This can be done two ways:
   - using the *root* account (simplest)
   - using a normal user account and then setting up *User Rights and Privileges*

   Either of these enables a suitable user who can administer the Linux environment as exposed to the MS Windows SMB/CIFS network environment.

   a) Configure root as the MS Windows administrator equivalent by running the following command and completing the required information as prompted:

```
# smbpasswd -a root

New SMB password: xxxxxxxxxx (does not have to be root password)
Retype new SMB password: xxxxxxxxx
Added user root.
```

   b) If you are using the root account, skip to step 5; if you are using a *normal user account* (not root), configure the local administrator account.

      NOTE: *You must set up an administrator account. Also, all user names, including administrator should be in lower-case, as Linux is case-sensitive.*

      i. Complete step a.

      NOTE: *This account will be removed or disabled after the administrator account has been established.*

      ii. Create a Linux account as follows:

```
# useradd -m -g 4 administrator
# passwd administrator
```

```
Enter new UNIX password: xxxxxxxxx
Retype new UNIX password; xxxxxxxx
Passwd: password updated successfully
```

iii. Add SMB/CIFS credentials as follows:

```
# smbpasswd -a administrator

New SMB password: xxxxxxxxx
Retype new SMB password: xxxxxxxxx
Added user administrator.
```

iv. Verify that the *administrator* account exists in the SMB/CIFS environment by running the following command and viewing its output:

```
# pdbedit -Lv administrator
Unix username:          administrator
NT username:
Home Directory:         \\CIFSFS\administrator
HomeDir Drive:
Logon Script:
Profile Path:           \\CIFSFS\administrator\profile
Domain:                 CIFSFS
Account desc:
Workstations:
Munged dial:
Logon time:             0
Logoff time:            9223372036854775807 seconds since the Epoch
Kickoff time:           9223372036854775807 seconds since the Epoch
Password last set:      Tue, 21 Sep 2010 09:30:00 CDT
Password can change:    Tue, 21 Sep 2010 09:30:00 CDT
Password must change:   never
Last bad password:      0
Bad password count:     0
Logon hours:            FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
```

v. Set administrator privileges for this account as shown here:

```
# net rpc rights grant "CIFSFS\administrator" SeMachineAccountPrivilege
SeTakeOwnershipPrivilege SeBackupPrivilege SeRestorePrivilege
SeRemoteShutdownPrivilege SePrintOperatorPrivilege SeAddUsersPrivilege
SeDiskOperatorPrivilege -Uroot%xxxxxxxxx
```

xxxxxx is the root password.

vi. Disable the *root* account by running the following command:

```
# smbpasswd -d root
Disabled user root.
```

vii. Delete the root account from the CIFS password back end by running the following command:

```
# pdbedit -u root -x
```

5. Add a local UNIX group, which is required for shared resource ownership and access control.

6. For each group identity (at least one is required) create a UNIX group and then map it into the SMB/CIFS environment as shown here:

```
# groupadd engineers
# net groupmap add unixgroup=engineers ntgroup=engineers type=local
```

   Replace *engineers* with an appropriate group name for the site.

7. Add local SMB/CIFS user accounts.

   a. For each separate user who requires read-only access to the SMB/CIFS server, add a UNIX used account and then create the SMB/CIFS account extensions as shown here:

```
# useradd -m -g users myname
# passwd myname
Enter new UNIX password: xxxxxxxxx
Retype new UNIX password: xxxxxxxxx
Passwd: password updated successfully

# smbpasswd -a myname
New SMB password: xxxxxxxxx
Retype new SMB password: xxxxxxxxx
Added user Administrator.
```

   b. For each separate user who requires read/write to access the SMB/CIFS server add a UNIX used account and then create the SMB/CIFS account extensions as shown here and provide the required information. The -G argument specifies that the user is assigned to a secondary group.

```
# useradd -m -g users myname
# usermod -a -G engineers myname
# passwd myname
Enter new UNIX password: xxxxxxxxx
Retype new UNIX password: xxxxxxxxx
Passwd: password updated successfully

# smbpasswd -a myname
New SMB password: xxxxxxxxx
Retype new SMB password: xxxxxxxxx
Added user Administrator.
```

   IMPORTANT: User accounts created in Linux and Unix are case-sensitive and case-preserving. Create accounts using only lower-case characters.

8. Configure the shared resources.

NEXT STEP: Configure Share Resources

Active Directory Domain Member Server

NOTE: *A customer's gateway can be configured only as standalone (local authentication) server (see* Stand-alone Server (Workgroup Authentication)*)OR as an Active Directory Domain member server. It cannot be configured as both; it must be one or the other.*

Active Directory configuration requires the following procedures.

- SMB configuration
- NTP configuration
- Edit the krb5.conf file
- Edit the nsswitch file
- Join the domain
- Validate that the domain has been joined

Microsoft Active Directory requires a fully functional DNS service to resolve machine names and identify critical services that enable or support Active Directory. The use of WINS (Windows Internetworking Name Server) is NOT necessary with Active Directory – in fact, larger sites mostly disable the use of NetBIOS over TCP/IP, thus nullifying the use of WINS.

DNS servers used with the CFS server should use the DNS server that is authoritative for the Active Directory domain when configured as an Active Directory domain member server.

NOTE: In the examples shown in this section, the following names are used:

- Active Directory domain controller = *w2k8r2.xyz.project.local*
- Realm name = *xyz.project.local*
- Windows machine name = *W2K8R2*
- *P*re-Windows 2000 domain name = *XYZ*

Edit the krb5 File

IMPORTANT: *Do NOT edit the /etc/krb5.conf file on a standalone server. This file must be edited only on an Active Directory domain member server.*

The */etc/krb5.conf* file must contain the correct domain name information for SMB to install successfully.

1. Run the following command to open the */etc/krb5.conf* file.

```
# vi /etc/krb5.conf
```

2. Edit the file as shown with your domain name information.

```
[libdefaults]
     default_realm =  XYZ.PROJECT.LOCAL

[realms]
     XYZ.PROJECT.LOCAL = {
```

```
        kdc   =  xxx.xxx.x.xx (AD DNS server; see Site Survey)
        kdc = w2k8r2.xyz.project.local
        admin_server   =  w2k8r2.xyz.project.local
    }
[domain_realm]
         .dxplatform.local = DXPLATFORM.LOCAL
        dxplatform.local = DXPLATFORM.LOCAL
```

NEXT STEP: Configure the SMB/CIFS Server

Configure the SMB/CIFS Server

BEFORE YOU BEGIN: Did you Edit the krb5 File?

1. Run the following command to open the */etc/resolv.conf* file.

```
#vi /etc/resolv.conf
```

2. Set the CFS DNS server address to the authoritative DNS server for the Active Directory domain as shown here:

```
domain xyz.domain.local
search xyz.domain.local other.dns.domain
nameserver xxx.xxx.x.xx
```

Where *xyz.domain.local* is the fully qualified DNS name for the Active Directory realm, other.dns.domain is any other domain, the address *xxx.xxx.x.xx* should be replaced with the correct IP address for the Active Directory DNS server.

NOTE: *The Network Time Protocol service on the CFS server should be configured using either the Microsoft Active Directory domain controller for the domain it will be joining, or using the same time server it has been set up to use. Dell recommends that you point the Domain Controller, CFS server and DX cluster to the same time source. To configure the Windows Time Service on the Domain Controller to use an external time source, refer to the KB article below:*

http://support.microsoft.com/kb/816042

3. Run the following command to save the original */etc/samba/smb.conf* file.

```
#mv /etc/samba/smb.conf /etc/samba/smb.conf.orig
```

4. Run the following command to open the */etc/samba/smb.conf* file.

```
#vi /etc/samba/smb.conf
```

5. In the file */etc/samba/smb.conf*, replace the *workgroup* name *AD,* the *realm* name, and the *netbios name* (in upper case characters – each max 14 characters) that is appropriate for the site:

```
[global]
        workgroup = XYZ
        realm = XYZ.PROJECT.LOCAL
        netbios name = CIFSFS
        server string = DX Storage
        security = ADS
        log level = 1
        log file = /var/log/samba/log.%L.%m
        max log size = 0
        smb ports = 445
        machine password timeout = 0
        load printers = No
        disable spoolss = Yes
        os level = 0
        ldap ssl = no
        idmap backend = tdb
        idmap uid = 5000000-10000000
        idmap gid = 5000000-10000000
        winbind separator = +
        winbind cache time = 3000
        winbind enum users = Yes
        winbind enum groups = Yes
        idmap config XYZ : backend = rid
        idmap config XYZ : range = 100000 - 2999999
        posix locking = No
```

NOTE: *A non-overlapping Idmap config entry should be added for each trusted domain that must access this server. The domain range must not clash or overlap with the Idmap UID and GID range and the Idmap config range specified in the example of the XYZ domain shown above.*

6. Join the domain as shown here:

```
# net ads join -Uadministrator%xxxxxxxxx (where xxxx is password)
Using short domain name -- XYZ
Joined 'CIFSFS' to realm 'xyz.project.local'
…
```

NOTE: SAMBA initiates DDNS update to register itself on DNS Server. If DNS Update fails, make sure the DNS server is setup to update dynamic updates. This configuration may vary depending on the Active directory Operating System. Refer to the following KB article for additional information: http://support.microsoft.com/kb/816592

7. Start the CFS SMB/CIFS server daemons as shown here:

```
# chkconfig winbind on
# chkconfig smb on
# service winbind start
# service smb start
```

8. Check the integrity of the domain trust account.

```
# wbinfo –t
checking the trust secret via RPC calls succeeded
```

9. Run the following command to obtain a list of Active Directory domain user accounts:

```
# wbinfo –u
XYZ+administrator
XYZ+guest
XYZ+krbtgt
XYZ+jthorely
XYZ+jackb
```

10. Run the following command to obtain the list of Active Directory domain group accounts:

```
# wbinfo –g
XYZ+domain computers
XYZ+domain controllers
XYZ+schema admins
XYZ+enterprise admins
XYZ+cert publishers
XYZ+domain admins
XYZ+domain users
XYZ+domain guests
XYZ+group policy creator owners
XYZ+ras and ias servers
XYZ+allowed rodc password replication group
XYZ+denied rodc password replication group
XYZ+read-only domain controllers
XYZ+enterprise read-only domain controllers
XYZ+dnsadmins
XYZ+dnsupdateproxy
```

NOTE: In a large Active Directory environment, the commands wbinfo -u and wbinfo -g may take a long time to complete.

11. Edit the file */etc/nsswitch.conf*.

Edit the entries:

```
passwd:     file
shadow:     file
group:      file
```

to the following:

```
passwd:     file winbind
shadow:     file winbind
group:      file winbind
```

12. Obtain a list of users via the NSS interface:

```
# getent passwd
root:x:0:0:root:/root:/bin/bash
daemon:x:1:1:daemon:/usr/sbin:/bin/sh
bin:x:2:2:bin:/bin:/bin/sh
sys:x:3:3:sys:/dev:/bin/sh
sync:x:4:65534:sync:/bin:/bin/sync
...
statd:x:116:65534::/var/lib/nfs:/bin/false
XYZ+administrator:*:5000000:5000000:Administrator:/home/XYZ/administrator:/bin/false
XYZ+guest:*:5000001:5000001:Guest:/home/XYZ/guest:/bin/false
XYZ+krbtgt:*:5000002:5000000:krbtgt:/home/XYZ/krbtgt:/bin/false
XYZ+jthorely:*:5000003:5000000:John H. Thorely:/home/XYZ/jthorely:/bin/false
XYZ+jackb:*:5000004:5000000:Jack B. Black:/home/XYZ/jackb:/bin/false
```

13. Obtain a list of groups:

```
# getent group
root:x:0:
daemon:x:1:
bin:x:2:
sys:x:3:
...
XYZ+domain computers:x:5000006:
XYZ+domain controllers:x:5000007:
XYZ+schema admins:x:5000008:XYZ+administrator
XYZ+enterprise admins:x:5000009:XYZ+administrator
XYZ+cert publishers:x:5000010:
XYZ+domain admins:x:5000011:XYZ+administrator
XYZ+domain users:x:5000000:
XYZ+domain guests:x:5000001:
XYZ+group policy creator owners:x:5000012:XYZ+administrator
XYZ+ras and ias servers:x:5000013:
XYZ+allowed rodc password replication group:x:5000014:
XYZ+denied rodc password replication group:x:5000015:XYZ+krbtgt
XYZ+read-only domain controllers:x:5000016:
XYZ+enterprise read-only domain controllers:x:5000017:
XYZ+dnsadmins:x:5000018:
XYZ+dnsupdateproxy:x:5000019:
XYZ+unixgroup:x:5000020:XYZ+jthorely
```

14. Validate the domain membership information:

```
# net ads info
LDAP server: 172.16.10.27
LDAP server name: w2k8r2.xyz.project.local
Realm: XYZ.PROJECT.LOCAL
Bind Path: dc=XYZ,dc=PROJECT,dc=LOCAL
LDAP port: 389
Server time: Mon, 20 Sep 2010 21:51:32 CDT
KDC server: 172.16.10.27
```

```
Server time offset: 0
```

15. Copy the contents of the file /var/lib/samba/smb_krb5/krb5.conf.<*domainname*> to the file /etc/krb5.conf.

    It should contain similar information as shown here:

```
[libdefaults]
        default_realm = XYZ.PROJECT.LOCAL
        default_tgs_enctypes = RC4-HMAC DES-CBC-CRC DES-CBC-MD5
        default_tkt_enctypes = RC4-HMAC DES-CBC-CRC DES-CBC-MD5
        preferred_enctypes = RC4-HMAC DES-CBC-CRC DES-CBC-MD5

[realms]
        XYZ.PROJECT.LOCAL = {
                kdc = 192.168.1.22
                kdc = w2k8r2.xyz.project.local
                admin_server = w2k8r2.xyz.project.locals

        }
```

16. The Domain Administrator or any other domain user must have access rights on the system to set permissions on the share. Check if the Domain Administrator has access rights on the system

```
# net rpc rights list accounts -UAdministrator%<password>

Domain Admins
SeMachineAccountPrivilege
SeTakeOwnershipPrivilege
SeBackupPrivilege
SeRestorePrivilege
SeRemoteShutdownPrivilege
SePrintOperatorPrivilege
SeAddUsersPrivilege
SeDiskOperatorPrivilege

BUILTIN\Print Operators
No privileges assigned

BUILTIN\Account Operators
No privileges assigned

BUILTIN\Backup Operators
No privileges assigned

BUILTIN\Server Operators
No privileges assigned
```

```
BUILTIN\Administrators
SeMachineAccountPrivilege
SeTakeOwnershipPrivilege
SeBackupPrivilege
SeRestorePrivilege
SeRemoteShutdownPrivilege
SePrintOperatorPrivilege
SeAddUsersPrivilege
SeDiskOperatorPrivilege

Everyone
No privileges assigned
```

17. If the Domain Admins group has no privileges assigned grant access rights to the Domain Admins group on the CFS Gateway system

```
#  net rpc rights grant "<Domain Name>\Domain Admins" SeMachineAccountPrivilege
SeTakeOwnershipPrivilege SeBackupPrivilege SeRestorePrivilege SeRemoteShutdownPrivilege
SePrintOperatorPrivilege SeAddUsersPrivilege SeDiskOperatorPrivilege -
Uadministrator%<passwd>
```

The CFS server is now ready for configuration of shared resources.

NEXT STEP: Configure Share Resources

NFS Gateway Service

The procedures in this section define how to configure services so shared resources can be added using the appropriate procedure detailed in the Shared Resource Configuration section of this document. The CFS file system resource can be accessed from a remote UNIX or Linux machine via the NFS version 3.

***NOTE**: Dell does not support NFS version 4 at this time.*

The following procedure configures NFS server only. After performing the procedure, you must then configure shared resources. See Configure Share Resources.

Before Configuring NFS

To configure NFS, you must ensure IP forwarding and the firewall are disabled.

Disable IP forwarding

By default, RedHat Linux has IP Forwarding disabled. To verify this, run the following command.

```
# sysctl net.ipv4.ip_forward
```

The command returns a value of 0, if IP forwarding is disabled, or 1 if it is enabled. IP forwarding should be disabled following the procedures in the RHEL6 documentation.

Configure NFS to share the CFS mount

NFS must be configured to share the CFS mount. The configuration installed includes NFS kernel support. Configure the NFS */etc/exports* file as shown in the Configure Share Resources section. Do not forget to turn the NFS service on and to start it. See the Configure Share Resources section for more information.

NEXT STEP: Configure Share Resources

Configure Share Resources

Share resources provide the connection points for the SMB/CIFS and NFS protocol gateway services that are provided by the CFS server. Each connection protocol has its own configuration requirements. These have been integrated into the *CFS-admin* utility.

SMB/CIFS Shared Resource Configuration

Specific configuration requirements of SMB/CIFS shared resources on a CFS server depend on whether the server is configured for local authentication or as a member server in an Microsoft Windows Active Directory security context (domain).

BEFORE YOU BEGIN: Did you set up Stand-alone Server (Workgroup Authentication)?

Add a Share (Standalone Server)

IMPORTANT: *Make sure the system is configured as Standalone Server. Do not continue with this procedure if the system is not a Standalone Server.*

1. Run the following command to open the */etc/samba/smb.conf* file.

```
#vi /etc/samba/smb.conf
```

2. Add a share stanza to the file as shown here:

```
[share_name]
        comment = Share1
        path = /mnt/share_name/toplevel
        read only = No
        use sendfile = Yes
```

NOTE: *Replace the* **share_name** *with an appropriate name. Dell recommends using the same name as was used to create the CFS mounted resource. The toplevel directory must be created within the CFS mount point because it creates the share-point for CIFS and NFS use.*

3. Create the *toplevel* directory.

```
# mkdir -p /mnt/share_name/toplevel
```

4. Set file system ownership and group ownership for the user and group that will have write access to the shared resource.

```
# chown -R myname:users /mnt/share_name/toplevel
# chmod -R ug+rw,o-rwx /mnt/share_name/toplevel
# find /mnt/share_name/toplevel -type d -exec chmod g+sx {} \;
```

In the example above, there is a CFS mounted file system resource under the mount point */mnt/share_name*. The contents of this directory will be owned by the user *gillian* and the group *Users*.

***NOTE**: Setting the SGID flag on directories enforces inheritance of group ownership of the top-level directory as new files and directories (folders) are created.*

***NOTE**: A read-only account can be a member of the same group that has write access to the share. The read-only user account can have write permissions in the file system; this account will be granted read-only access through ACL settings on the share itself. This is completed in the following stp.*

5. Using a Microsoft Windows workstation or server (XP or later), set the share ACL.
   a) Open a CMD terminal session. Run the following:

```
C:\Users\Administrator> net use * /d
…

Do you want to continue this operation? <Y/N> [N]: y
The command completed successfully.
```

If you are asked to continue the operation, answer **Yes** (see above).

This disconnects all open connections to remote systems. This is important in order to avoid the side effect of open connections that can impede the ability to access security objects on the remote system.

   b) Launch Windows Explorer and type \\NetBIOS name of the CFS server (for example, \\CIFSFS) and press <Enter>.
      Windows Security will display for authentication.
      i. Type in the NetBIOS name of the server and the administrator password (for example CIFSFS\administrator)
      ii. In the **Password** field, enter the password that you created for the administrator account on the CFS server.
      iii. Press <Enter>.

   After a few moments, the shares on the CFS server should display.

   c) Click the Start button, type MMC in the search box, and press <Enter> to launch the console.
   d) Click on **File** and in select **Add/Remove Snap-in**.
   e) From the left panel (Available snap-ins), select **Computer Management**.
   f) Click the **Add** button.

g) Click the button to select **Another computer**.
h) In the field provided **Browse** to the CFS machine, or enter the NetBIOS name of the CFS server, and then click **Finish**.
i) Click **OK**.
j) Click (+) to expand the **Computer Management** tree.
k) Click (+) to expand the **System Tools** tree.
l) Click (+) to expand **Shared Folders**.
m) Click **Shares** to see the shares that are available.
n) Double-click the share on which access controls must be set.
o) In the **Properties** dialog, click the **Share Permissions** tab.
p) Click **Add**.
q) In the **Select Groups/Users** dialog, click **Advanced**.
r) Click **Find Now**.
s) Select a group that should have access, and click **OK**. (Do this for as many groups that require access to this share.)
t) Click **OK** again.
u) Set the access permissions required for the group **Everyone**. (If you do not wish to allow all users access, do **NOT** set Deny permissions, as this will lock every user out. Instead, delete the group Everyone from the Access Control List.)
v) Click **Apply**.
Permissions are now set on the share. Anyone who is not a member of the group cannot connect to the share.

w) Click **OK**.
x) Close the Microsoft Management Console and return to the CMD terminal.
y) Run the following command to disconnect all open connections to the CFS server.

```
C:\Users\Administrator> net use * /d
…

Do you want to continue this operation? <Y/N> [N]: y
The command completed successfully.
```

Add a share (Active Directory Domain Member)

BEFORE YOU BEGIN: Did you set up the Active Directory Domain Member Server and set up the DX6000G Cluster File Server?

IMPORTANT: *Make sure that the system is configured as an Active Directory Domain Member Server. Do not continue with this procedure if the system is not an Active Directory Domain Member Server.*

1. Run the following command to open the */etc/samba/smb.conf* file:

```
#vi /etc/samba/smb.conf
```

2. Add a share stanza to the file as shown here:

```
[share_name]
```

```
        comment = ShareName
        path = /mnt/share_name/toplevel
        read only = No
        use sendfile = Yes
```

NOTE: *Replace the **share_name** with an appropriate name. Dell recommends using the same name that was used to create the CFS mounted resource.*

3. Set file system ownership and group ownership for the user and group that will have write access to the shared resource.

   In the following example there is a CFS mounted file system resource under the mount point */mnt/share_name/toplevel*.

   NOTE: The *toplevel* directory must be created.

4. The contents of this directory will be owned by the user *Gillian* and the group *Users*. Set ownership and permissions as shown here:

```
# chown -R XYZ+gillian:XYZ+"domain users" /mnt/share_name/toplevel
# chmod -R ug+rw,o-rwx /mnt/ share_name/toplevel
# find /mnt/ share_name/toplevel -type d -exec chmod g+sx {} \;
```

   Any Active Directory user or group name that has a space in it must be within quotation marks, as shown in this example.

   NOTE: *Setting the SGID flag on directories enforces inheritance of group ownership of the top-level directory as new files and directories (folders) get created.*

5. Some sites require enforced Active Directory–based ACL (Access Control List) inheritance (See Site-Survey responses to determine whether this is the case). If this occurs, add the following to the share (e.g.: *[share_name]*) stanza of the */etc/samba/smb.conf* configuration file:

```
acl group control = Yes
force unknown acl user = Yes
inherit acls = Yes
inherit owner = Yes
inherit permissions = Yes
map acl inherit = Yes
```

   NOTE: *These settings CANNOT be overridden from a Microsoft Windows client, even if the user attempting to make the change is a Domain Administrator.*

Remove a Share

To remove a share, simply delete the share stanza and all its parametric contents. Open the */etc/samba/smb.conf* file with an editor, locate the share stanza, and delete the stanza and all contents down to the first blank line.

NOTE: *You do not need to restart the smbd daemon (or any others) when a share stanza is removed.*

NFS Shared Resource Configuration

Configuration requirements for NFS shared resources on a CFS server are affected by the NFS version or versions that must be supported. The following procedures step through the configuration issues that must be taken into account.

For NFS version 3, the specified fsid is optional and can be any 32-bit number and must be unique among all the exported file systems. For NFS version 4, the fsid for the root of the NFSv4 export tree must be 0.

Add an NFSv3 Share Point

1. Add the nfs mount specification - edit the */etc/exports* file:

   The following is a sample entry, using a mount point with the name "CFS1" and specifying a "rw" option for read/write access and a "root_squash" option to prevent root write access:

   ```
   /mnt/share_name/toplevel *(rw,root_squash)
   ```

   For greater security, specify which clients can access the exported share as shown in the following example:

   ```
   /mnt/share_name/toplevel 192.168.50.0/24(rw,root_squash)
   ```

2. Start NFS using the following command as root:

   ```
   # service nfs start
   # chkconfig nfs on
   ```

Remove an NFS Share Point

To remove an NFS mount resource, simply comment out or remove the entry from the */etc/exports* file.

NOTE: After every change to this file, the NFS Server must be restarted. Changes are not dynamically picked up as they are for a CIFS shared resource.

Upgrading CFS

The DX6000G Cluster File Server can be upgraded from its previous versions 2.5 and 2.6.x with minimal interruption. Administrators should plan for 15-20 minutes of downtime during the upgrade process.

1. Extract the latest version of the Dell-DX-CFS Software.zip as root.

   ```
   # unzip Dell-DX-CFS Software.zip
   ```

2. In the extracted folder, copy the latest CFS SW to root

```
# cp dell-dxcfs-2.x.x-x86_64.zip /root/
```

3. Unzip the CFS software

```
# unzip dell-dxcfs-2.x.x-x86_64.zip
```

Follow instructions in Section 2.4 Upgrading CFS in the CFS Installation Setup and Config Guide to upgrade your CFS software

Administrative Maintenance Procedures

Starting CFS and CNS

The Content Name Space should always be started prior to CFS.

1. Start Content Name Space with the following command:

```
# service caringo-cns start
```

   If any of several critical configuration parameters is missing or invalid, CNS will fail to start and will display an error message. When the configuration is corrected in the cns-admin script, CNS should start correctly.

2. Boot the CFS server.

   CFS will start automatically if the "mount on boot" option was selected during the configuration process for each mount point. If the process was stopped for any reason it can be manually started with a standard mount command.

   To mount all configured CFS mount points at once, run the following command:

```
# service caringo-cfs start
```

   To mount a single mount point previously defined using the *CFS-admin mkfs* script, run a command similar to the following (where */mnt/CFS1* is the mounted destination for the desired mount point)

```
# mount /mnt/CFS1
```

   If any of several critical configuration parameters is missing or invalid, CFS will fail to mount a share and will display an error message. Once the configuration is corrected in the *CFS-admin mkfs* script, the share should mount correctly.

3. Before writing any data to a mount point, run a *mount* command with no options to ensure the desired mount points started correctly and are present in the mounted list.

   If the mount point is not in the mounted list, the install was not successful and you should not write any data to the mount point.

Shut Down CFS and CNS

For the CFS server to cleanly shut down, SMB/CIFS and NFS services must be stopped. If these services were configured following the procedures outlined in the Gateway Protocol Configuration section of this document, the services will automatically stop in the correct order as the system is shut down. If the services do not automatically stop, you will need to manually shut them down.

To manually stop the SMB/CIFS and NFS services, use the following commands:

```
# service smb stop
# service nmb stop
# service winbind stop
# service nfs stop
```

NOTE: *If any of these services are already stopped, the above command may show as failed for that service.*

To stop CFS and un-mount all configured shares, use the following command:

```
# service caringo-cfs stop
```

To stop and/or un-mount a specific configured share, use a command similar to the following (where */mnt/CFS1* is the mounted destination for the mount point):

```
# umount /mnt/CFS1
```

NOTE: *If CFS is stopped using a kill command, a fusermount -u /mnt/mount_point command must be executed before restarting to ensure the mount point is properly released and remounted. If the remount option is utilized, the mount point will be un-mounted and then immediately mounted. After CFS has been stopped, CNS may also be stopped using the following command:*

```
# service caringo-cns stop
```

Special Considerations for MD3200i Spooler

When a CFS node using MD3200i external spooler storage is being shut down, the following procedure must be used to avoid a possible hang during shutdown. If these steps are not followed, the system may hang and need to be powered down or reset manually.

1. Unmount any spool directories mounted from the MD3200i.

```
# umount /var/cache/cns
# umount /var/spool/cfs/<spool1>

And so forth for any other spools
```

2. Log out iscsi sessions.

```
# iscsiadm –m node –u
```

3. Flush multipath cache:

```
# multipath –F
```

4. Shut down or restart the system.

Additional Information

See the *DX Storage Cluster File Server (CFS) Setup and Configuration Guide* for information on the following topics:

- File Revisions and DX Object Storage Deletion
- Timescapes
- DX Object Storage Metadata and Policies
- Temp and Logging Space Consideration

Appendix A. Gateway Protocol Support

This appendix provides information that may be useful to the CFS Protocol Gateway server administrator or implementer.

Protocol Gateway Limitations

The CFS Protocol Gateway administrator or implementer should note the Dell DX Object Storage cluster provides support for a metadata-rich set of attributes that may be used to describe binary objects that are being stored. The use of file sharing protocols limits how these may be used since attributes that are not known to the underlying network file system topography cannot be utilized.

Supported Protocols

CFS can be used with any file sharing technology however, Dell's development efforts to date have focused primarily on SMB/CIFS and NFS. The focus of this deployment guide is upon specific support for SMB/CIFS and NFS.

Access Control Lists

POSIX ACL (Access Control List) metadata will be mapped into the Dell DX Object Storage HTTP SCSP metadata header content only if the underlying file system has been mounted with POSIX compliant ACL support and with Extended Attributes (EAs) enabled.

Where the CFS protocol gateway is used to access the Dell DX Object Storage via the SMB/CIFS protocols, Microsoft Windows NTFS and NTFS5 ACLs will be mapped to the nearest equivalent POSIX ACLs, but this will be possible only where the underlying file system has been mounted with support for POSIX ACLs and Extended Attributes.

It should be noted the mapping of MS Windows NTFS ACLs will be affected not only by a closely approximated mapping to POSIX ACLs, but additionally may be overridden by specific share specification parameters that can be used to enforce access controls in such manner that even the MS Windows network administrator cannot change or override them. Description of such controls is beyond the scope of this document.

SMB/CIFS Protocol Support

The CFS SMB/CIFS server makes use of Samba version 3.4.7 (or later). This application fully implements all documented SMB (Server Message Block) and Microsoft Windows CIFS (Common Internet File System) protocols. Samba has a complete implementation of these protocols however, the behavior of Samba as the SMB/CIFS server in a Microsoft Windows network environment is determined by settings in the */etc/samba/smb.conf* file. For example, Samba can be configured so only certain SMB protocols are supported by setting the value of the *max protocol* parameter.

The current default value of the *max protocol* parameter is *NT1* (the latest CIFS support level). Samba version 3.5.0 (and above) also supports the new Microsoft Windows Vista (and above) *SMB2* protocol. This protocol will be enabled by default in Samba version 3.6.0.

Samba can be configured to support the following SMB/CIFS protocols:

CORE, COREPLUS, LANMAN1, LANMAN2, NT1, SMB2

Sites that elect to use a *max protocol* setting other than default do so at their own discretion outside of Dell's supported configurations.

NOTE: Other Samba configuration parameters can be set in the *[global]* stanza, or in a share stanza, that can impact connection protocol behavior. Dell recommends operation of the CFS protocol gateway server only within Dell supported boundaries.

Appendix B. NFS Client Guidelines

***NOTE**: Configuration requirements for NFS shared resources on a NFS client are dependent on the operating system platform used. Adjust all commands appropriately according to the operating system vendor's documentation.*

If not already present, install the required software packages for the NFS client machine by running the following command as root:

```
# yum install portmap nfs
```

Mount an NFS share on the NFS client machine, where the <sharename> matches the ones specified in the /etc/exports file:

```
# mount –t nfs \
-o hard,intr,bg,tcp,rsize=1048576,wsize=1048576,nordirplus \
host_name:/share_name mount-dir'
```

Using the maximum *rsize* and *wsize* values is highly recommended, as they can greatly improve network transfer speeds for CFS by ensuring the largest possible block size is always transmitted. The *nordirplus* is also recommended to improve directory listing performance. Please see the nfs(5) man page [http://www.rt.com/man/nfs.5.html] for additional performance tuning parameters for NFS clients.

***NOTE:** Mounting NFS from an OS X server requires a manual mount command similar to the above, as CFS requires several non-standard options that cannot be set via Finder. Specifically, OS X NFS mounts*

for CFS require the 'nolock' option to function correctly. Also, if mounting as a non-root user on an OS X client, users will either need to add the "insecure" option on the NFS server to allow the server to accept packets sent from a non-privileged (> 1024) port or mount as root using "sudo" and add the "resvport" option to the mount options. For additional OS X specific parameters for NFS clients please see the OS X mount_nfs(8) man page at:
[http://developer.apple.com/DOCUMENTATION/Darwin/Reference/ManPages/man8/mount_nfs.8.html]

Appendix C Manual Configuration Procedures

Much of the DX Object File Storage Gateway configuration is automated. This section provides information on how to manually configure those parts. It is intended for informational purposes only so that you can understand what is happening during the process.

NOTE: *These procedures are provided primarily as information about the activities that occur during the installation scripts. Perform these procedures only if there is a problem with the installation scripts.*

Create the Master Boot Record (MBR) on the Second Drive

NOTE: *This procedure is only valid for single-server configurations.*

By creating a master boot record on a second logical volume, you can always ensure that the system is bootable if a boot drive is removed, or – in the less predictable event – that the boot order is randomly swapped by the operating system.

1. Open a terminal session and run the following command to determine which drive is the boot drive.

```
# df | grep boot
```

2. One of the following will display:

```
/dev/sda1   198337   30414   157683   17%   /boot
```

OR

```
/dev/sdb1   198337   30414   157683   17%   /boot
```

This is the device and partition the boot drive is on.

3. Run the following command to start grub.

```
# grub
```

4. Based on the information obtained about the boot drive, set up the master boot record. (This example assumes that the command in step one discovered drive 0 as the boot drive.)

If the boot drive is sda, enter the following:

```
# grub> root (hd0,0)

Filesystem type is ext2fs, partition type 0x83
```

If the boot drive is sdb, enter the following:

```
# grub> root (hd1,0)

Filesystem type is ext2fs, partition type 0x83
```

5. Set up the first master boot record

```
# grub> setup (hd0)

Checking if "/boot/grub/stage1" exists... no
Checking if "/grub/stage1" exists... yes
Checking if "/grub/stage2" exists... yes
Checking if "/grub/e2fs_stage1_5" exists... yes
Running "embed /grub/e2fs_stage1_5 (hd0)"...  26 sectors are embedded.
succeeded
Running "install /grub/stage1 (hd0) (hd0)1+26 p (hd0,0)/grub/stage2
/grub/grub.conf"... succeeded
Done.
```

6. Set up the second master boot record

```
# grub> setup (hd1)

setup (hd1)
 Checking if "/boot/grub/stage1" exists... no
 Checking if "/grub/stage1" exists... yes
 Checking if "/grub/stage2" exists... yes
 Checking if "/grub/e2fs_stage1_5" exists... yes
 Running "embed /grub/e2fs_stage1_5 (hd1)"...  26 sectors are embedded.
succeeded
 Running "install /grub/stage1 d (hd1) (hd1)1+26 p (hd0,0)/grub/stage2
/grub/grub.conf"... succeeded
Done.
```

Disable SELinux

After installing Red Enterprise Linux, SELinux is disabled.

1. Open the configuration file.

```
# cd /etc/selinux/
# vi config
```

2. Change the settings as follows to disable SELinux.

```
SELINUX=disabled
```

3. Reboot the system.

   NOTE: *Do not continue to configure the system until it has been rebooted.*

Create the YUM Repository and Install Packages

The CFS installation process is dependent on additional rpm packages that are not installed on the system by default. These packages are available on the Red Hat Enterprise Linux distribution media included with the system. Running these packages requires a local YUM repository.

To create a local YUM repository on your system:

1. Ensure the CFS is powered on.
2. Insert the operating system media that came with the system into the optical drive and allow the file system to auto mount.

   The default directory path for the auto mounted file system is */media/RHELx.x\x86_64\ DVD*. The white spaces in this file path cause errors during the YUM setup process.

3. Create an .iso image of RHEL6.

```
# dd if=/dev/dvd of=RHEL-6.0_x86_64.iso
```

4. Create the yum repository.

```
# mkdir /root/RHEL6
# mount -o loop,ro /root/RHEL-6.0-x86_64.iso /root/RHEL6
```

   NOTE: *Do not use root as your yum repository. Create a designated folder for the repository (such as RHEL6 in the example above).*

   NOTE: *The mount must be recreated each time you reboot the system.*

5. Remove any cached packages from the system and enable the local YUM repository.

```
# yum clean all
# yum repolist
```

6. Edit the repository to remove packagekit-media.repo.

```
# cd /etc/yum.repos.d
# rm packagekit-media.repo
# vi rhel6.repo
```

   NOTE: *The packagekit-media.repo file must be deleted each time you reboot the system.*

7. Add the following to *rhel6.repo* file.

```
[InstallMedia]
name=Red Hat Enterprise Linux 6.0
mediaid=1285193176.460470
metadata_expire=-1
gpgcheck=0
cost=500
baseurl=file:/root/RHEL6
enabled=1
```

8. Ensure that avahi-daemon is running.

```
# service avahi-daemon status
```

If avahi-daemon is not running, use the following command:

```
# service avahi-daemon start
# chkconfig avahi-daemon on
```

9. Install the packages required to complete CFS installation.

a. Create a package list file.

```
# vi pkglist
```

b. Add the following services to the file.

```
dialog
libtdb
libldb
libtalloc
samba
samba-common
samba-client
samba-winbind
samba-winbind-clients
tdb-tools
ntpd
iscsi-initiator-utils
device-mapper-multipath
device-mapper-multipath-libs
```

NOTE: *Check the spelling of each of the above entries carefully. Also, when installing the packages, if any dependencies are identified, go ahead and install it now and repeat the installation of any package that failed because of the dependency.*

10. Run the following command:

```
# for pkg in `cat pkglist`
do
 yum install $pkg --nogpgcheck
```

```
done
```

The CFS system is now ready to be updated with all the dependencies required to complete the installation.

Stop and Disable Services

1. Create a list of file services.

```
# vi list
```

2. Add the following services to the file.

```
Nscd
NetworkManager
nmb
smb
winbind
ntpd
nfs
iptables
ip6tables
```

3. Run the following command:

```
# for svc in `cat list`
do
 service $svc stop
 chkconfig $svc off
done
```

***NOTE**: If any of these services are already stopped, the above command may show as failed for that service.*

Set up the NTP Server

The CFS server(s) must use the same NTP time source as the domain controllers that will be used for handling Active Directory-based credentials. Even when Active Directory is not used, it is still recommended to use a common time source for all CFS servers.

***NOTE**: If you selected an NTP server while installing the operating system on the CFS, you do not need to perform the following procedure.*

1. Verify that ntpd is not running.

```
# service ntpd status
```

2. If NTP is running, stop the service.

```
# service ntpd stop
```

3. Edit the /etc/ntp.conf file.

```
# vi /etc/ntp.conf
```

4. Edit the file to configure the time server to your site time server as identified in the site survey form.

```
server clock.xyz.project.local stratum 2
server time1.nis.gov stratum 1
```

IMPORTANT: Use only appropriate entries. If using external NTP servers, make sure you are authorized to use those servers.

***NOTE**: The clock should be the time server of your Windows domain controller server and is shared between the Windows domain controller and the CFS server. It is very important that the time be set within 5 seconds of datum (Atomic Clock time).*

5. Restart the time service and configure autostart on reboot.

```
# chkconfig ntpd on
# service ntpd start
```

Configure the Network Interfaces for Bonding

Channel bonding enables two or more network interfaces to act as one, simultaneously increasing the bandwidth and providing redundancy. Dell recommends the following bonds for the networks that are part of the CFS solution.

	Bond	Ethernet ports	Network
Single-Server Solution	0	0-1	CFS/NFS Gateway Network
	1	2-5	DX Cluster Network
Failover Solution	0	0-1	CFS/NFS Gateway Network
	1	2-5	DX Cluster Network and iSCSI storage

***NOTE**: If iSCSI is on a separate network, use the following port designations.*

	Bond	Ethernet ports	Network
Single-Server Solution	0	0-1	CFS/NFS Gateway Network
	1	2-5	DX Cluster Network
Failover Solution	0	0-1	CFS/NFS Gateway Network
	1	2-3	DX Cluster Network

	2	4-5	iSCSI Storage Network

Dell supports two different types of bonding: balance-alb (adaptive load balancing) and link aggregation control (LACP, also known as 802.3ad). Balance-alb is configured as mode=6; 802.3ad is configured as mode=4. (see below). You should deploy the type of bonding that the customer site is most comfortable with.

Configuration of Ethernet bonding under RHEL 6.0 requires the configuration of bond master files (one per bond), and a configuration file for each of its slave ports.

***NOTE**: This procedure requires extensive information about the customer's network. You should have the completed Site Survey form readily available.*

1. Change to the network-scripts directory.

```
# cd /etc/sysconfig/network-scripts
```

2. Create interface configuration files (one required for each bonded network).

```
# vi ifcfg-bondn  (where n is the bond number, beginning with 0)
```

3. Enter the following information in the file, replacing the network addresses with those used in your network.

```
DEVICE=bond0
ONBOOT=yes
IPADDR=xxx.xx.x.xx  (see Site Survey)
BOOTPROTO=none
PREFIX=24  (appropriate netmask significant bits from the Site Survey)
IPV6INIT=no
NAME="System bond0"
TYPE=Ethernet
GATEWAY=xx.xx.x.x  (see Site Survey)
DEFROUTE=yes
IPV4_FAILURE_FATAL=yes
BONDING_OPTS="mode=6 miimon=100"  (see Site Survey)
DNS1= xxx.xxx.xxx.xxx  (see Site Survey).
DOMAIN=xyz.project.local  (see Site Survey)
```

***NOTE**: DNS and Gateway information may not be required for the bond created for the public network.*
***NOTE**: Red Hat may alter the PREFIX value. Verify that this value matches what is shown in ifconfig.*

The next step defines the configuration for each of the Ethernet ports that are to be bonded.

4. Create a configuration file for the Ethernet port.

```
# vi ifcfg-ethn  (where n is the NIC port number, starting from 0)
```

NOTE: *The NIC port number and its MAC address can be obtained from /etc/udev/rules.d/70-persistent-net.rules. This file identifies all of the NIC ports found when the system was last powered on.*

5. Enter the following information.

```
DEVICE="ethn"  (where n is the number of the Ethernet port)
ONBOOT=yes
HWADDR=00:26:B9:3D:55:19  (validate this against /etc/udev/rules.d/70-persistent-net.rules)
NAME="System ethn"  (where n is the number of the Ethernet port)
BOOTPROTO=none
MASTER=bondx  (where x is the number of the bond to which the Ethernet port belongs)
SLAVE=yes
USERCTL=no
```

Repeat the above steps for every Ethernet port, substituting the appropriate port number (e.g., eth1 port, replacing eth0 with eth1).

6. Repeat steps 2-5 for the remaining bond(s) and assigned Ethernet ports.

7. Load the kernel module to validate the channel bonding interfaces.

   a. As root user, go to the */etc/modprobe.d* directory.

   b. Create a ***bonding.conf*** file.

   c. In the ***bonding.conf*** file, insert the following lines:

```
alias bond0 bonding
alias bond1 bonding
```

NOTE: *Add one entry for each bonded Ethernet interface that has been configured.*

8. Restart network services

```
# service network restart
```

9. Reboot the system.

Configure the Domain Name Service Resolver

A CFS Gateway requires an authoritative domain name server. Before beginning this procedure, obtain the following information from Windows Active Directory Services:

- Domain name
- Domain name server IP address
- Open the */etc/resolv.conf* file.

```
# vi /etc/resolv.conf
```

1. Enter the following information in the *resolv.conf* file:

```
search xyz.project.local (see Site Survey)
nameserver xx.xx.x.x (see Site Survey)
domainname xyz.project.local (see Site Survey)
```

***NOTE**: Make sure there is a DNS entry for the CIFS/NFS interface of the server.*

2. Edit the */etc/nsswitch.conf* file.

```
# vi /etc/nsswitch.conf
```

3. Between the *shadow:* and *networks:* lines, you will find the following:

```
…
hosts: files dns
…-
```

Change this line to:

```
…
hosts: files dns mdns4_minimal [NOTFOUND=return] dns mdns4
…-
```

4. Restart network services.

```
# service network restart
```

Install the CFS Software

The *caringo-cfs* package is available as a Red Hat rpm package that is installed with a shell script. As the root user, install the package and dependencies with the following command.

1. Verify that all services are stopped.

```
# service nscd status
# service smb status
# service nmb status
# service winbind status
# service nfs status
```

***NOTE**: Services should show as not running for all of the above. If they are running, they must be disabled.*

***NOTE**: If any of these commands fail (other than saying that the service is not running), this indicates that the required software package was not installed. See Create the YUM Repository and Install Packages. Make sure the RHEL6 iso is still mounted. If it is not mounted, run the following command:*

```
# mount -o loop,ro /root/RHEL-6.0-x86_64.iso /root/RHEL6
```

1. Copy the CFS installation zip file to /root and extract it.

2. Change directory to the newly extracted directory tree.

3. Change the execute permissions.

```
# chmod 755 installDXCFS.sh
```

4. Install the CFS package.

```
# ./installDXCFS.sh
```

Spooler and Cache File Systems

The spooler is a shared file system that serves as a spool/cache for files before they are written to DX Object Storage. The spooler also contains journals and file revision information. Depending on the solution, the spooler can be on the CFS server itself (single-server solution) or on an external storage device (failover solution).

As the spooler grows, unnecessary spooler entries are eventually evicted. Also, if the files are needed for later reads, they are accessed directly from the spooler, if the files are still on the spooler.

NOTE: *Each CFS share must have its own dedicated spooler partition and its own cache eviction/management process.*

If the spooler partition is shared between CFS shares, there will be eviction conflicts between the processes. Because of this restriction, some customers mount subfolders inside a single CFS share via CIFS. You need to make a new CFS share for any of the following scenarios:

- When you need to specify different lifepoint policies or custom metadata for the mount. For example, PACS data has a retention period of 15 years, whereas emails are kept for 3 years.

- When applications have different usage patterns (lots of small files vs lots of very large file transfers). The spooler has a max files configuration set by default to 100,000. Eviction is based on a partition's used capacity as well as max files. This prevents a 100Gb spooler partition from filling up with 50M small files.

- When you want to isolate the performance impact of your applications from each other. For example, a customer running a batch job that must quickly scan all files in a share for viruses, would quickly flood the cache and evict files from another application running in the same share (affecting read performance).
- If the CNS cache must be installed on its own spooler file system.

Create disk partitions and spooler file systems (Single-server solution)

A local spooler on the CFS should have a capacity of 2TB (for SATA drives) or 600GB (for SAS drives) and be configured as RAID 1.

1. Find the drive that the spooler will be residing on. Substitute for Drive_ID below the correct drive ID. (Could be sda or sdb; run the #df to identify the disks.)

2. Dedicate a drive for the spooler.

```
# fdisk /dev/Drive_ID
```

Create a partition that spans the whole drive:

a. Type n and press <Enter> (to create a new drive)

b. Select the partition type by typing p (primary) and press <Enter>

c. Enter Partition number = **1**.

d. For First cylinder use default **1**.

e. Press <Enter> at ending partition = end of disk.

f. Type w (to write the partition).

3. Make /dev/sda1 an LVM physical disk.

```
# pvcreate /dev/sda1
```

4. Create a volume group on sda1.

```
# vgcreate CacheVG /dev/sda1
```

5. Review the volume group

```
# vgdisplay CacheVG
```

- **Free PE** — The first number (before the slash) is the number of physical extents (PEs) the disk is divided into.
- **PE size** (a few lines up) shows how big each PE is.

6. Create a Logical Volume (LV) for the CNS cache.

```
# lvcreate -l <PEcount> -n cns_cache CacheVG
```

<PEcount> is the number of PEs to use for CNS cache.

7. Create LVs for each CFS spool.

```
# lvcreate -l <PEcount> -n <CFSname>_spool CacheVG
```

<PEcount> can be different for each CFS spool you are creating (and different for what was used for CNS).

8. For each CFS spool, run the following command.

```
# mkfs -t ext4 /dev/mapper/CacheVG-<fsname>
```

<fsname> is 'cns_cache', '<CFSname>_spool', etc.

9. Create a directory for each of the file systems that will be mounted.

```
# mkdir -p /var/spool/cfs/share_name
```

10. Create a directory for the CNS spool cache.

```
# mkdir -p /var/cache/cns
```

11. Edit the /etc/fstab file to add entries for the each of the spoolers you created and the cache.

```
# vi /etc/fstab
```

12. Add the following after the last line.

```
/dev/mapper/<CFSname>_spool   /var/spool/cfs/share_name ext4  rw,acl,user_xattr,nodelalloc 0 0
/dev/mapper/CacheVG-cns_cache   /var/cache/cns ext4   rw,acl,user_xattr,nodelalloc 0 0
```

13. Mount the file system.

```
# mount -a
```

14. Validate that the mount occurred.

```
# df | grep cfs
```

External Spooler File System (Failover solution)

The preferred external storage option (documented in this guide) is the MD3200i. However, many installations may already have an external storage infrastructure in place. The CFS gateway can use any external storage solution that meets the following criteria:

- Access to the external storage is supported on RHEL 6; this includes considerations of performance, availability, etc.
- External storage supports creation of an ext3 or ext4 file system.
- File system can be mounted (non-concurrently) on each gateway system.
- External storage system is highly reliable (write operations performed to the external storage and returned as completed must actually be completed, they cannot be lost in e.g. power failure, network outage, etc).
- Connection to the external storage can be done either via the existing Ethernet adapter configuration required for the gateway (in the case of e.g. iSCSI-based storage) or via another connection that does not interfere with any connectivity required for the gateway (e.g. Fibre Channel, SAS, etc). Ethernet-based storage may be located on any of the networks connected to the gateway; choice of a network should take into consideration bandwidth requirements, network addressing, etc.

NOTE: Configuration of non-MD3200i external storage is beyond the scope of this document; the documentation for the external storage solution should be used for any required configuration.

Configure the MD3200i

The MD3200i spooler supports a minimum of six drives and a maximum of 20. The drives should be configured as RAID5 with one hot-spare. You will also need to create a number of LUNs (logical disks). The number of LUNs you should create is based on the number of CFS shares, maximum file size, and performance requirements. You will need one LUN for the CNS cache and one for each CFS file system.

1. Cable the MD3200i as follows:

    a. Connect the management ports on each controller to the public network.

    b. Connect the data ports (4 on each controller) to the storage network.

2. Install the PowerVault Modular Disk Storage software (also known as MDCU) on a Windows or Linux management station.

    IMPORTANT: *Do not install any Dell PowerVault Linux drivers.*

    NOTE: *The management station should be separate from the CFS or CSN servers.*

3. Run the management software and allow it to automatically discover storage devices.

    NOTE: *Autodiscovery assumes that the management station is on the same subnet of the public network. If the Autodiscovery does not begin automatically when launching the application, select Tools→Autodiscovery in the console.*

4. Configure the disk group.

    a. Open the MD3200i management screen by double-clicking on storage array.

    b. Click the **Logical** tab, right-click a disk, and click **Create** to configure the drives on the MD3200i into a single RAID5 array, leaving one disk as hotspare.

    c. Click the **Logical** tab, right-click an array, and click **Create** to configure the LUNs.

        NOTE: *Use essentially the same logic as you did for sizing the storage nodes on a DX cluster when creating virtual disks in the shared spool. For example, if the customer has larger file sizes, you should create larger LUNs. If you have reason to believe that the customer may add additional CFS file systems in the future, you can create extra LUNs for them to use when necessary, or leave unconfigured space for expansion.*

5. Create a host group on the MD3200i for each failover node pair.
    a. Mappings menu → Define → Host Group
    b. Enter host group name and click OK.

6. Assign the LUNs you created to the host group with whatever LUN numbers are desired.

    a. Expand Undefined Mappings, right click on the LUN and then select Define Additional Mappings.

    b. Select Host group or host.

    c. Select LUN # and click Add.

7. Configure iSCSI on the MD3200i.

    a. Setup tab.

b. Configure iSCSI Host Ports
c. IP address and subnet mask
d. Don't need gateway
e. Select iSCSI host ports for other data ports and configure the same.
f. If VLAN, click Advanced IPv4 Settings and enter VLAN information.
g. Click OK.
h. Check Manage iSCSI Settings and Target Authentication set to None.

8. From a root login on the CFS node, ping all eight iSCSI IPs to ensure they are working.

9. Change the initiator name.

```
# iscsi-iname -p iqn.2010-12.local.project.xyz >> /etc/iscsi/initiatorname.iscsi
# vi /etc/iscsi/initiatorname.iscsi
```

(Replace *local.project.xyz* with the IANA-qualified name of the site's iscsi device.)

When you open the file, you will see the following.

```
InitiatorName= iqn.1994-05.com.redhat:8072808d1ba6
iqn.2010-12.local.project.xyz:a82389cde897
```

10. Move the cursor to the first "I" of iqn (in the InitiatorName line), delete to the end of the line (Shift+d), join the two lines (Shift+j), and delete the space between the = sign and the newly generated iqn.

11. Start iscsid and iscsi.

```
# service iscsid restart
# service iscsi restart
```

NOTE: *These services will not start if the iqn name is not set correctly. Recommended naming convention is iqn.year-month.system-name-hexnumber.*

12. Set iscsi and iscsid to start on boot

```
# chkconfig iscsid on
# chkconfig iscsi on
# chkconfig multipath on
```

13. Discover the iSCSI ports.

```
# iscsiadm -m discovery -t st -p 172.16.16.30
```

The listing should show all iSCSI ports on the MD3200i.

14. Open the MD storage console to make the LUNs visible to the host:
    a. Mappings tab.
    b. Select the host group with the LUNs, and select **Define→Host**.

c. Enter a **User label**.
d. Select **Add by selecting a known unassociated host port identifier**.
e. Select an entry from the **known unassociated host port identifier** drop-down list.
NOTE: Click the Refresh button if a host port does not initially appear in the list.
f. Enter a host name and Click **Add**, and click **Next**
g. Select **Linux** and then **Finish**.

This allows the host to see the LUNs; prior to this step all it could see was an 'access volume', which is used for in-band management of the storage.

15. Retrieve the LUN information.

```
# iscsiadm -m discovery -t st -p 172.16.16.30
```

16. Run the following command to log the server into the storage, and create disks in */dev* and mapper entries for the multipath disk-mapper volumes.

```
# iscsiadm -m node -l
```

(-l is the letter "l")

NOTE: *If the multipath.conf file is not created by default, copy it from /usr/share/doc/device-mapper-multipath-0.4.9/multipath.conf to /etc/ and reload the multipath daemon.*

```
# service multipathd reload
```

17. Display all LUNs in the host group.

```
# multipath -ll
```

All the LUNs in the host group should be displayed (Use size information to verify the correct LUNs are displayed).

18. Run the following command to create a /dev/mapper entry for the new partition. This command will need to be repeated for each LUN being used.

```
# kpartx -a /dev/mapper/mpath<d>
```

19. Run the following command to create a partition table where *mpath<d>* (e.g. mpathe) is one of the LUN names that was shown when the multipath -ll command was run. This command will need to be repeated for each LUN being used. Options used for the fdisk command should be the same as those used below in step 2 of the single-server solution.

```
# fdisk /dev/mapper/mpath<d>
```

a. Create a partition that spans the whole drive:

i. Type n and press <Enter> (to create a new drive)

ii. Select the partition type by typing p (primary) and press <Enter>

iii. Enter Partition number = **1**.

iv. For First cylinder use default **1**.

v. Press <Enter> at ending partition = end of disk.

vi. Type w (to write the partition).

20. Create an ext4 file system on the partition.

    This command will need to be repeated for each LUN being used. The mapper entry will generally be of the form mpath<d>p1 (e.g. mpathep1).

```
# mkfs -t ext4 /dev/mapper/mpath<d>p1
```

21. Add the LUNs to /etc/fstab. A number of lines will be added, of the form:

```
/dev/mapper/mpath<d>p1 /var/spool/cfs/cifs1   ext4   acl,user_xattr,nodelalloc,_netdev 0 0
```

    The first entry should be for /var/cache/cns instead of /var/spool/cfs/cifs1. The name 'cifs1' should be chosen to correspond to the customer's intended CFS names.

22. After all LUNs are added to /etc/fstab, mount the new file systems.

```
# mount -a
```

23. Use 'df' to verify that all are mounted.

    WARNING: *Do **NOT** mount the drives on the backup node while the file system is mounted on the primary node. This will cause the mount to fail and could even damage the file system. You must unmount the file system from the primary drive BEFORE mounting the file system on the backup node. Likewise, even after assigning a mount on the backup node (assuming you have unmounted from the primary node), do **NOT** set it to automount.*

24. Verify that */etc/iscsi/iscsid.conf* has `node.startup = automatic` enabled.

```
# iscsiadm -m node -o show  | egrep 'node\.(name|conn\[0\]\.startup|startup)'
```

    Expected output is as follows:

```
node.name = iqn.2010-12.local.project.xyz.cifs1
node.startup = automatic
node.conn[0].startup = manual
node.name = iqn.2010-12.local.project.xyz.cifs1
node.startup = automatic
node.conn[0].startup = manual
```

25. Set the iSCSI initiator to startup automatically at boot-time and login to the iSCSI targets at start-up.

```
# iscsiadm -m node -o update -n 'node.conn[0].startup' -v automatic
```

26. Validate that the new settings have been accepted by running following command (previously used in the step above):

```
# iscsiadm -m node -o show  | egrep 'node\.(name|conn\[0\]\.startup|startup)'
```

Expected output is as follows:

```
node.name = iqn.2010-12.local.project.xyz.cifs1
node.startup = automatic
node.conn[0].startup = automatic
node.name = iqn.2010-12.local.project.xyz.cifs1
node.startup = automatic
node.conn[0].startup = automatic
```

27. Restart the iscsi initiator service, and set it to run automatically by executing the following commands.

```
# service iscsi stop
# service iscsid stop
```

```
# chkconfig iscsi  --level 2345 on && chkconfig iscsi --list
# chkconfig iscsid --level 2345 on && chkconfig iscsid --list
```

```
# service iscsid start
# service iscsi start
# service multipathd start
```

28. Edit the Caringo startup scripts to make sure that CNS and CFS start in the proper order.

    a. In the file */etc/init.d/caringo-cns*, change the following:

```
# Required-Start:    $networking $avahi $time
# Required-Stop:     $networking $avahi $time
```

    To

```
# Required-Start:    $networking $avahi $time $remote_fs
# Required-Stop:     $networking $avahi $time $remote_fs
```

    b. In the file /etc/init.d/caringo-cfs; change the following:

```
# Required-Start:    $networking $avahi $time
# Required-Stop:     $networking $avahi $time
```

    To

```
# Required-Start:    $networking $avahi $time caringo-cns
# Required-Stop:     $networking $avahi $time caringo-cns
```

c. Run the following commands:

```
# chkconfig --del caringo-cns
# chkconfig --del caringo-cfs
# chkconfig --add caringo-cns
# chkconfig --add caringo-cfs
```

29. Repeat these steps on the backup node.

IMPORTANT: *Of the MD3200i configuration steps, only the one adding a host to the host group need to be repeated. Do* **NOT** *repeat the fdisk and mkfs steps.*

Configure the Cluster Name Space (CNS)

See the *DX Storage Cluster File Server (CFS) Setup and Configuration Guide*.

Configure the CFS and its DX Object Storage Mount Points

See the *DX Storage Cluster File Server (CFS) Setup and Configuration Guide*.